



Installing and Configuring Windows Server® 2012



Exam Ref 70-410

Craig Zacker

Sample Chapters

Copyright © 2012 by Craig Zacker

All rights reserved.

To learn more about this book visit:

<http://go.microsoft.com/fwlink/?Linkid=272594>

Contents

Introduction	xi
<i>Microsoft certifications</i>	<i>xi</i>
<i>Errata & book support</i>	<i>xii</i>
<i>We want to hear from you</i>	<i>xii</i>
<i>Stay in touch</i>	<i>xii</i>
<i>Preparing for the exam</i>	<i>xiii</i>
Chapter 1 Installing and configuring servers	1
Objective 1.1: Install servers	2
Planning for a server installation	2
Choosing installation options	6
Upgrading servers	12
Migrating roles	14
Objective summary	16
Objective review	17
Objective 1.2: Configure servers	18
Completing postinstallation tasks	18
Using Server Manager	26
Configuring services	36
Delegating server administration	37
Objective summary	38
Objective review	39
Objective 1.3: Configure local storage	40
Planning server storage	40
Understanding Windows disk settings	42

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Working with disks	45
Objective summary	62
Objective review	63
Answers.....	66
Chapter 2 Configure server roles and features	71
Objective 2.1: Configure file and share access	71
Creating folder shares	72
Assigning permissions	77
Configuring Volume Shadow Copies	86
Configuring NTFS quotas	87
Objective summary	88
Objective review	89
Objective 2.2: Configure print and document services	91
Deploying a print server	91
Sharing a printer	97
Managing documents	101
Managing printers	102
Using the Print and Document Services role	104
Objective summary	109
Objective review	109
Objective 2.3: Configure servers for remote management	111
Using Server Manager for remote management	112
Using Remote Server Administration Tools	119
Working with remote servers	120
Objective summary	120
Objective review	121
Answers.....	123
Chapter 3 Configure Hyper-V	129
Objective 3.1: Create and configure virtual machine settings	129
Virtualization architectures	130
Hyper-V implementations	131
Installing Hyper-V	134

Using Hyper-V Manager	136
Configuring resource metering	148
Objective summary	149
Objective review	149
Objective 3.2: Create and configure virtual machine storage.	151
Virtual disk formats	152
Creating virtual disks	153
Configuring pass-through disks	159
Modifying virtual disks	160
Creating snapshots	161
Connecting to a SAN	162
Objective summary	167
Objective review	168
Objective 3.3: Create and configure virtual networks	169
Creating virtual switches	170
Creating virtual network adapters	176
Creating virtual network configurations	180
Objective summary	181
Objective review	182
Answers.	184

Chapter 4 Deploying and configuring core network services 189

Objective 4.1: Configure IPv4 and IPv6 addressing	189
IPv4 addressing	190
IPv6 addressing	197
Planning an IP transition	201
Objective summary	205
Objective review	205
Objective 4.2: Configure servers	207
Understanding DHCP	207
Deploying a DHCP server	214
Deploying a DHCP relay agent	219
Objective summary	222
Objective review	222

Objective 4.3: Deploy and configure the DNS service	223
Understanding the DNS architecture	224
Deploying a DNS server	233
Objective summary	240
Objective review	241
Answers.	243
Chapter 5 Install and administer Active Directory	249
Objective 5.1: Install domain controllers	249
Deploying Active Directory Domain Services	250
Objective summary	264
Objective review	265
Objective 5.2: Create and manage Active Directory users and computers	267
Creating user objects	267
Creating computer objects	277
Managing Active Directory objects	280
Objective summary	285
Objective review	285
Objective 5.3: Create and manage Active Directory groups and organizational units (OUs)	287
Working with groups	292
Objective summary	300
Objective review	301
Answers.	303
Chapter 6 Create and manage Group Policy	307
Objective 6.1: Create Group Policy objects (GPOs).	307
Understanding Group Policy objects	308
Configuring a Central Store	309
Using the Group Policy Management console	309
Managing starter GPOs	312
Configuring Group Policy settings	313
Creating multiple local GPOs	314

Objective summary	316
Objective review	316
Objective 6.2: Configure security policies	317
Defining local policies	318
Using security templates	322
Configuring local users and groups	325
Configuring User Account Control	329
Objective summary	332
Objective review	332
Objective 6.3: Configure application restriction policies.....	334
Using software restriction policies	334
Using AppLocker	341
Objective summary	344
Objective review	344
Objective 6.4: Configure Windows Firewall	346
Understanding Windows Firewall settings	346
Working with Windows Firewall	347
Using the Windows Firewall control panel	348
Using the Windows Firewall with Advanced Security console	352
Objective summary	357
Objective review	357
Answers.....	360
 <i>Index</i>	 367

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Configure server roles and features

This chapter covers some of the fundamental services that most Windows servers perform. In the business world, file and printer sharing were the reasons computers were networked in the first place, and with Windows Server 2012, remote management has become a critical element of server administration.

Objectives in this chapter:

- Objective 2.1: Configure file and share access.
- Objective 2.2: Configure print and document services.
- Objective 2.3: Configure servers for remote management.

Objective 2.1: Configure file and share access

One of the critical daily functions of server administrators is deciding where users should store their files and who should be permitted to access them.

This objective covers how to:

- Create and configure shares
- Configure share permissions
- Configure offline files
- Configure NTFS permissions
- Configure access-based enumeration (ABE)
- Configure Volume Shadow Copy Service (VSS)
- Configure NTFS quotas

Creating folder shares

Sharing folders makes them accessible to network users. After you have configured the disks on a file server, you must create shares to enable network users to access those disks. As noted in the planning discussions in Chapter 1, “Installing and configuring servers,” you should have a sharing strategy in place by the time you are ready to create your shares. This strategy should consist of the following information:

- What folders you will share
- What names you will assign to the shares
- What permissions you will grant users to the shares
- What Offline Files settings you will use for the shares

If you are the Creator Owner of a folder, you can share it on a Windows Server 2012 computer by right-clicking the folder in any File Explorer window, selecting Share With > Specific People from the shortcut menu, and following the instructions in the File Sharing dialog box, as shown in Figure 2-1.

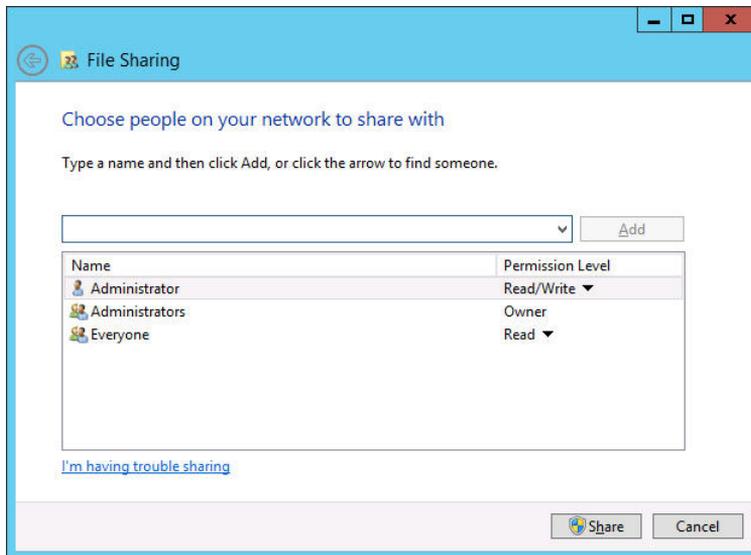


FIGURE 2-1 The File Sharing dialog box.

This method of creating shares provides a simplified interface that contains only limited control over elements such as share permissions. You can specify only that the share users receive Read or Read/Write permissions to the share. If you are not the Creator Owner of the folder, you can access the Sharing tab of the folder's Properties sheet instead. Clicking the Share button launches the same dialog box, and the Advanced Sharing button displays the Advanced Sharing dialog box, shown in Figure 2-2, which provides greater control over share permissions.

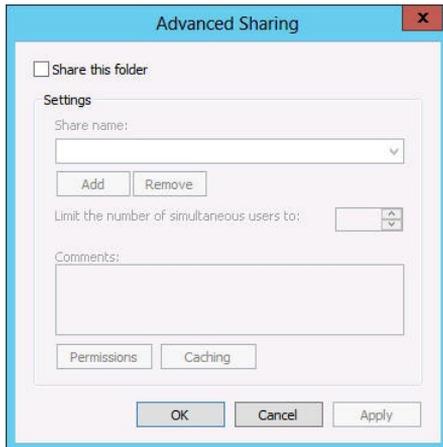


FIGURE 2-2 The Advanced Sharing dialog box.

NOTE NETWORK DISCOVERY

For the users on the network to be able to see the shares you create on the file server, you must make sure the Network Discovery and File Sharing settings are turned on in the Network and Sharing Center control panel.

To take control of the shares on all your disks on all your servers and exercise granular control over their properties, you can use the File and Storage Services home page in Server Manager.

Windows Server 2012 supports two types of folder shares:

- **Server Message Blocks (SMB)** SMB is the standard file sharing protocol used by all versions of Windows.
- **Network File System (NFS)** NFS is the standard file sharing protocol used by most UNIX and Linux distributions.

When you install Windows Server 2012, the setup program installs the Storage Services role service in the File and Storage Services role by default. However, before you can create and manage SMB shares by using Server Manager, you must install the File Server role service, and to create NFS shares, you must install the Server for NFS role service.

To create a folder share by using Server Manager, use the following procedure.

1. Log on to Windows Server 2012 using an account with Administrator privileges. The Server Manager window opens.
2. Click the File and Storage Services icon and, in the submenu that appears, click Shares to open the Shares home page.
3. From the Tasks menu, select New Share. The New Share Wizard starts, displaying the Select The Profile For This Share page, as shown in Figure 2-3.

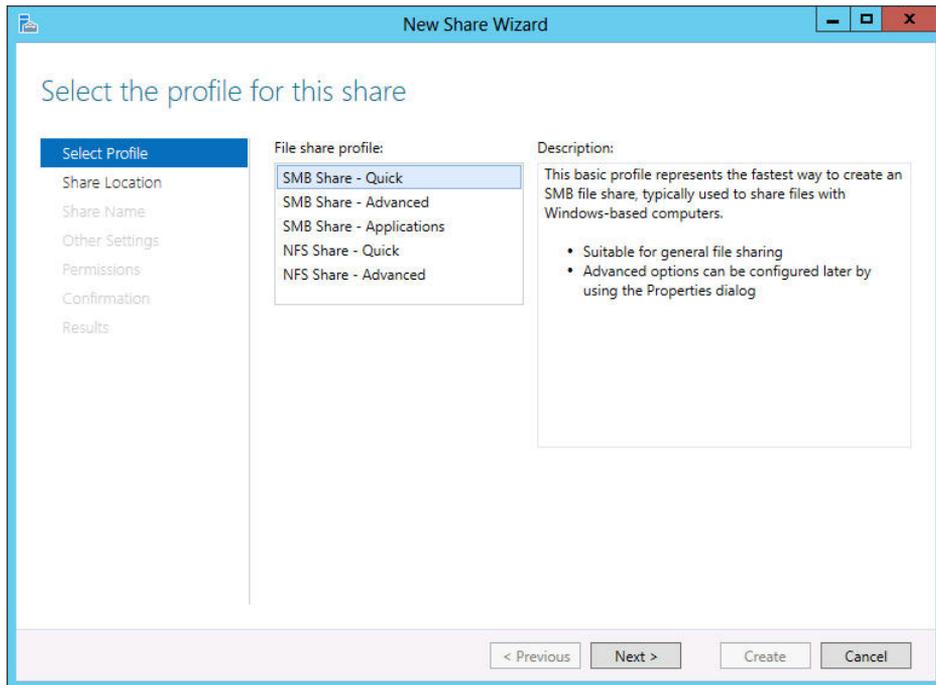


FIGURE 2-3 The Select The Profile For This Share page in the New Share Wizard.

4. From the File Share Profile list, select one of the following options:
 - **SMB Share–Quick** Provides basic SMB sharing with full share and NTFS permissions
 - **SMB Share–Advanced** Provides SMB sharing with full share and NTFS permissions and access to services provided by File Server Resource Manager
 - **SMB Share–Applications** Provides SMB sharing with settings suitable for Hyper-V and other applications
 - **NFS Share–Quick** Provides basic NFS sharing with authentication and permissions
 - **NFS Share–Advanced** Provides NFS sharing with authentication and permissions and access to services provided by File Server Resource Manager
5. Click Next. The Select The Server And Path For This Share page appears.
6. Select the server on which you want to create the share and either select a volume on the server or specify a path to the folder you want to share. Click Next. The Specify Share Name page appears.

MORE INFO NFS SHARING

Selecting one of the NFS share profiles adds two pages to the wizard: Specify Authentication Methods and Specify The Share Permissions. Both these pages provide access to functions implemented by the Server for NFS role service, as covered in Objective 2.1, “Configure Advanced File Services,” in Exam 70-412, “Configuring Advanced Windows Server 2012 Services.”

7. In the Share Name text box, specify the name you want to assign to the share and click Next. The Configure Share Settings page appears, as shown in Figure 2-4.

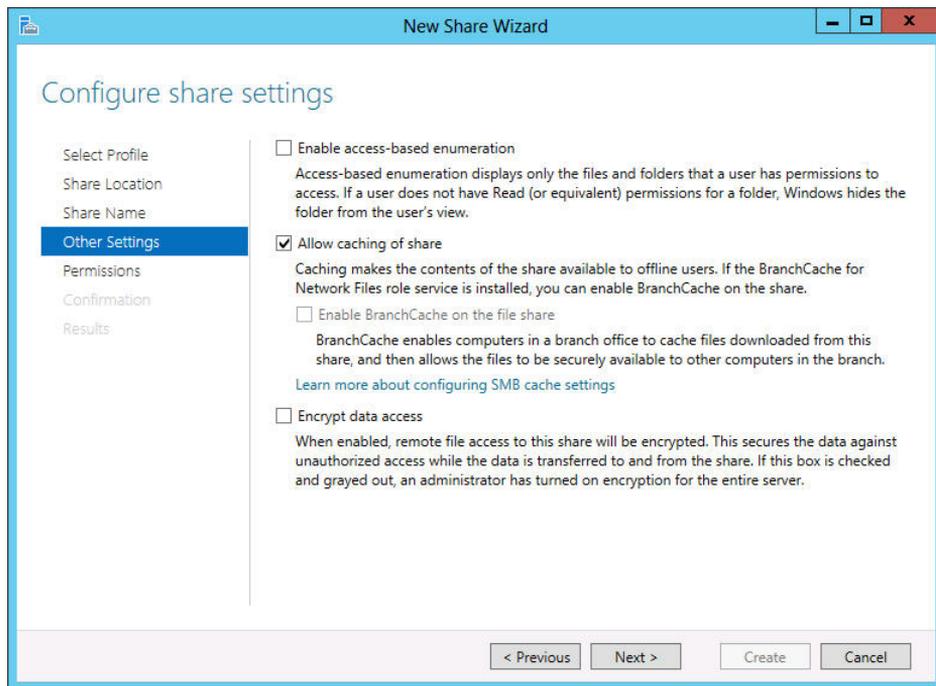


FIGURE 2-4 The Configure Share Settings page of the New Share Wizard.

8. Select any or all of the following options:
 - **Enable Access-Based Enumeration** Prevents users from seeing files and folders they do not have permission to access
 - **Allow Caching Of Share** Enables offline users to access the contents of this share
 - **Enable BranchCache On The File Share** Enables BranchCache servers to cache files accessed from this share
 - **Encrypt Data Access** Causes the server to encrypt remote file access to this share

NOTE ACCESS-BASED ENUMERATION

Access-based enumeration (ABE), a feature first introduced in Windows Server 2003 R2, applies filters to shared folders based on the individual user's permissions to the files and subfolders in the share. Simply put, users who cannot access a particular shared resource are unable to see that resource on the network. This feature prevents users from searching through files and folders they cannot access. You can enable or disable ABE for a share at any time by opening the share's Properties sheet in the Sharing and Storage Management console and clicking Advanced, which displays the same Advanced dialog box displayed by the Provision a Shared Folder Wizard.

NOTE OFFLINE FILES

Offline Files, also known as client-side caching, is a Windows feature that enables client systems to maintain local copies of files they access from server shares. When a client selects the Always Available Offline option for a server-based file, folder, or share, the client system copies the selected data to the local drive and updates it regularly so the client user can always access it, even if the server is offline. To enable clients to use the Offline Files feature, the share must have the Allow Caching Of Share check box selected. Windows Server 2012 and Windows 8 also have a new Always Offline mode for the Offline Files feature that causes clients to always use the cached copy of server files, providing better performance. To implement this mode, you must set the Configure slow-link mode Group Policy setting on the client to a value of 1 millisecond.

9. Click Next to move to the Specify Permissions To Control Access page.
10. Modify the default share and NTFS permissions as needed and click Next. The Confirm Selections page appears.

NOTE ADVANCED SHARE PROFILES

Selecting one of the Advanced share profiles adds two pages to the wizard: Specify Folder Management Properties and Apply A Quota To A Folder Or Volume. Both these pages provide access to functions of the File Server Resource Manager application, as covered in Objective 2.2, "Configure File Server Resource Manager (FSRM)," in Exam 70-411, "Administering Windows Server 2012."

11. Click Create. The View Results page appears as the wizard creates the share.
12. Close the New Share Wizard.

After you create a share by using the wizard, the new share appears in the Shares tile on the Shares home page in Server Manager. You can now use the tile to manage a share by right-clicking it and opening its Properties sheet or by clicking Stop Sharing.

Assigning permissions

Earlier in this chapter, you learned about controlling access to a file server to provide network users the access they need while protecting other files against possible intrusion and damage, whether deliberate or not. To implement this access control, Windows Server 2012 uses permissions.

Permissions are privileges granted to specific system entities, such as users, groups, or computers, enabling them to perform a task or access a resource. For example, you can grant a specific user permission to read a file while denying that same user the permissions needed to modify or delete the file.

Windows Server 2012 has several sets of permissions, which operate independently of each other. For the purpose of file sharing, you should be familiar with the operation of the following permission systems:

- **Share permissions** Control access to folders over a network. To access a file over a network, a user must have appropriate share permissions (and appropriate NTFS permissions if the shared folder is on an NTFS volume).
- **NTFS permissions** Control access to the files and folders stored on disk volumes formatted with the NTFS file system. To access a file, either on the local system or over a network, a user must have the appropriate NTFS permissions.

All these permission systems operate independently of each other and sometimes combine to provide increased protection to a specific resource. For network users to be able to access a shared folder on an NTFS drive, you must grant them both share permissions and NTFS permissions. As you saw earlier, you can grant these permissions as part of the share creation process, but you can also modify the permissions at any time afterward.

Understanding the Windows permission architecture

To store the permissions, each of these elements has an access control list (ACL). An ACL is a collection of individual permissions in the form of access control entries (ACEs). Each ACE consists of a security principal (that is, the name of the user, group, or computer granted the permissions) and the specific permissions assigned to that security principal. When you manage permissions in any of the Windows Server 2012 permission systems, you are actually creating and modifying the ACEs in an ACL.

To manage permissions in Windows Server 2012, you can use a tab in the protected element's Properties sheet, like the one shown in Figure 2-5, with the security principals listed at the top and the permissions associated with them at the bottom. Share permissions are typically found on a Share Permissions tab, and NTFS permissions are located on a Security tab. All the Windows permission systems use the same basic interface, although the permissions themselves differ. Server Manager also provides access to NTFS and share permissions by using a slightly different interface.

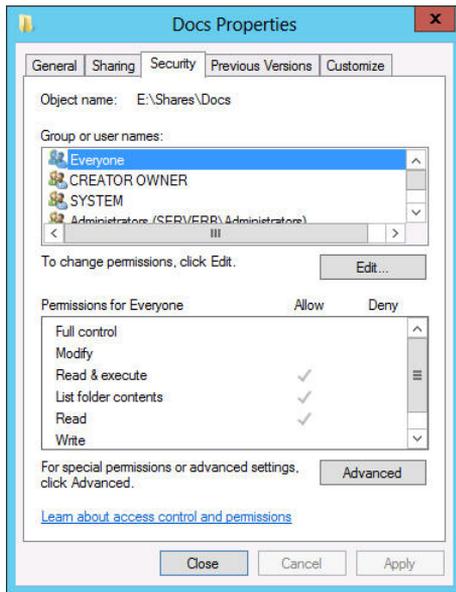


FIGURE 2-5 The Security tab of a Properties dialog box.

Understanding basic and advanced permissions

The permissions protecting a particular system element are not like the keys to a lock, which provide either full access or no access at all. Permissions are designed to be granular, enabling you to grant specific degrees of access to security principals.

To provide this granularity, each of the Windows permission systems has an assortment of permissions you can assign to a security principal in any combination. Depending on the permission system with which you are working, you might have dozens of different permissions available for a single system element.

Windows provides preconfigured permission combinations suitable for most common access control chores. When you open the Properties sheet for a system element and look at its Security tab, the NTFS permissions you see are called basic permissions. Basic permissions are actually combinations of advanced permissions, which provide the most granular control over the element.



EXAM TIP

Prior to Windows Server 2012, basic permissions were known as standard permissions and advanced permissions were known as special permissions. Candidates for certification exams should be aware of these alternative terms.

For example, the NTFS permission system has 14 advanced permissions you can assign to a folder or file. However, there are also six basic permissions, which are various combinations of

the 14 advanced permissions. In most cases, administrators work only with basic permissions. Many administrators rarely, if ever, work directly with advanced permissions.

If you find it necessary to work directly with advanced permissions, Windows makes it possible. When you click the Advanced button on the Security tab of any Properties sheet, an Advanced Security Settings dialog box appears, as shown in Figure 2-6, which enables you to access directly the ACEs for the selected system element. System Manager provides access to the same dialog box through a share's Properties sheet.

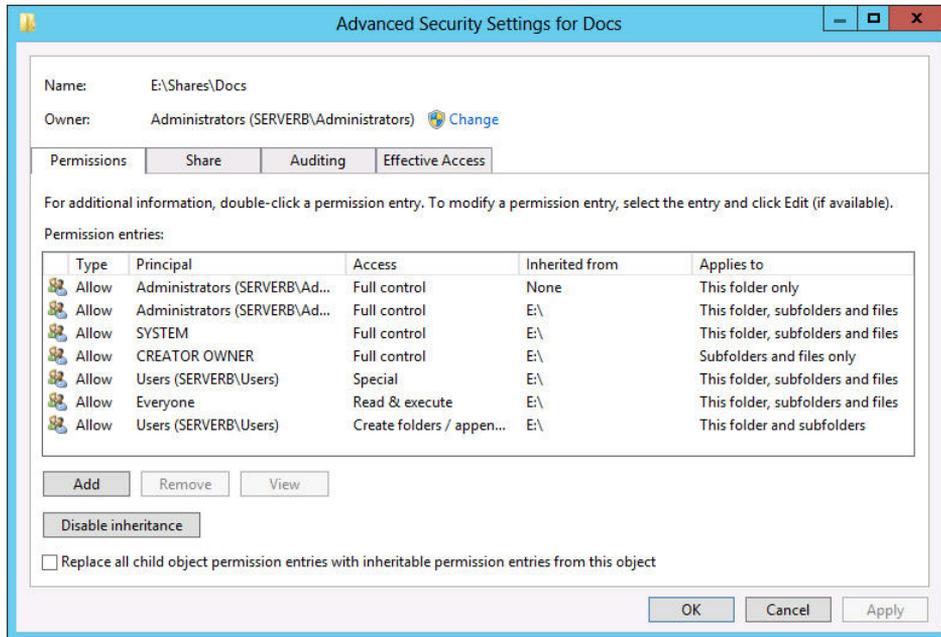


FIGURE 2-6 The Advanced Security Settings dialog box.

Allowing and denying permissions

When you assign permissions to a system element, you are, in effect, creating a new ACE in the element's ACL. There are two basic types of ACE: Allow and Deny. This makes it possible to approach permission management tasks from two directions:

- **Additive** Start with no permissions and then grant Allow permissions to individual security principals to give them the access they need.
- **Subtractive** Start by granting all possible Allow permissions to individual security principals, giving them full control over the system element, and then grant them Deny permissions for the access you don't want them to have.

Most administrators prefer the additive approach, because Windows, by default, attempts to limit access to important system elements. In a properly designed permission hierarchy, the use of Deny permissions is often unnecessary. Many administrators frown on their use,

because combining Allow and Deny permissions in a hierarchy can make it difficult to determine the effective permissions for a specific system element.

Inheriting permissions

The most important principle in permission management is that permissions tend to run downward through a hierarchy. This is called permission inheritance. Permission inheritance means that parent elements pass their permissions down to their subordinate elements. For example, when you grant Alice Allow permissions to access the root of the D drive, all the folders and subfolders on the D drive inherit those permissions, and Alice can access them.

The principle of inheritance greatly simplifies the permission assignment process. Without it, you would have to grant security principals individual Allow permissions for every file, folder, share, object, and key they need to access. With inheritance, you can grant access to an entire file system by creating one set of Allow permissions.

In most cases, whether consciously or not, system administrators take inheritance into account when they design their file systems and Active Directory Domain Services trees. The location of a system element in a hierarchy is often based on how the administrators plan to assign permissions.

In some situations, an administrator might want to prevent subordinate elements from inheriting permissions from their parents. There are two ways to do this:

- **Turn off inheritance** When you assign advanced permissions, you can configure an ACE not to pass its permissions down to its subordinate elements. This effectively blocks the inheritance process.
- **Deny permissions** When you assign a Deny permission to a system element, it overrides any Allow permissions that the element might have inherited from its parent objects.

Understanding effective access

A security principal can receive permissions in many ways, and it is important for an administrator to understand how these permissions interact. The combination of Allow permissions and Deny permissions a security principal receives for a given system element, whether explicitly assigned, inherited, or received through a group membership, is called the effective access for that element. Because a security principal can receive permissions from so many sources, it is not unusual for those permissions to conflict. The following rules define how the permissions combine to form the effective access.

- **Allow permissions are cumulative.** When a security principal receives Allow permissions from more than one source, the permissions are combined to form the effective access permissions.
- **Deny permissions override Allow permissions.** When a security principal receives Allow permissions, whether explicitly, by inheritance, or from a group, you can override those permissions by granting the principal Deny permissions of the same type.

- **Explicit permissions take precedence over inherited permissions.** When a security principal receives permissions by inheriting them from a parent or from group memberships, you can override those permissions by explicitly assigning contradicting permissions to the security principal itself.

Of course, instead of examining and evaluating all the possible permission sources, you can just open the Advanced Security Settings dialog box and click the Effective Access tab. On this tab, you can select a user, group, or device and view its effective access, with or without the influence provided by specific groups.

Setting share permissions

On Windows Server 2012, shared folders have their own permission system, which is independent from the other Windows permission systems. For network users to access shares on a file server, you must grant them the appropriate share permissions. By default, the Everyone special identity receives the Allow Full Control share permission to any new shares you create.

To modify the share permissions for an existing share by using File Explorer, you open the Properties sheet for the shared folder, select the Sharing tab, and then click Advanced Sharing and Permissions to open the Share Permissions tab, as shown in Figure 2-7.

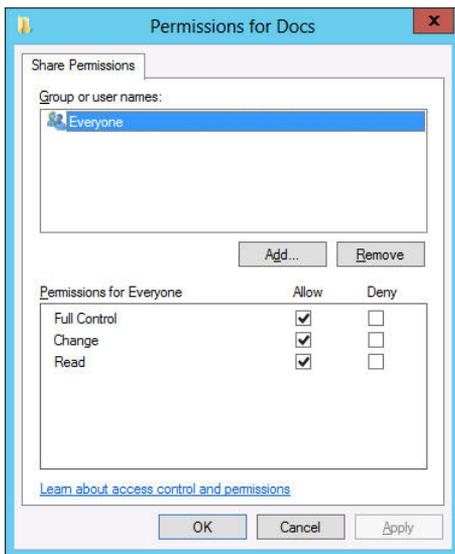


FIGURE 2-7 The Share Permissions tab for a shared folder.

By using this interface, you can add security principals and allow or deny them the three share permissions. To set share permissions by using Server Manager, either while creating a share or modifying an existing one, use the following procedure.

1. Log on to Windows Server 2012 and launch Server Manager.
2. Click the File and Storage Services icon and, in the submenu that appears, click Shares to open the Shares home page.

3. In the Shares tile, right-click a share and, from the shortcut menu, select Properties. The Properties sheet for the share opens.
4. Click Permissions. The Permissions page opens.
5. Click Customize Permissions. The Advanced Security Settings dialog box for the share opens.
6. Click the Share tab to display the interface shown in Figure 2-8.

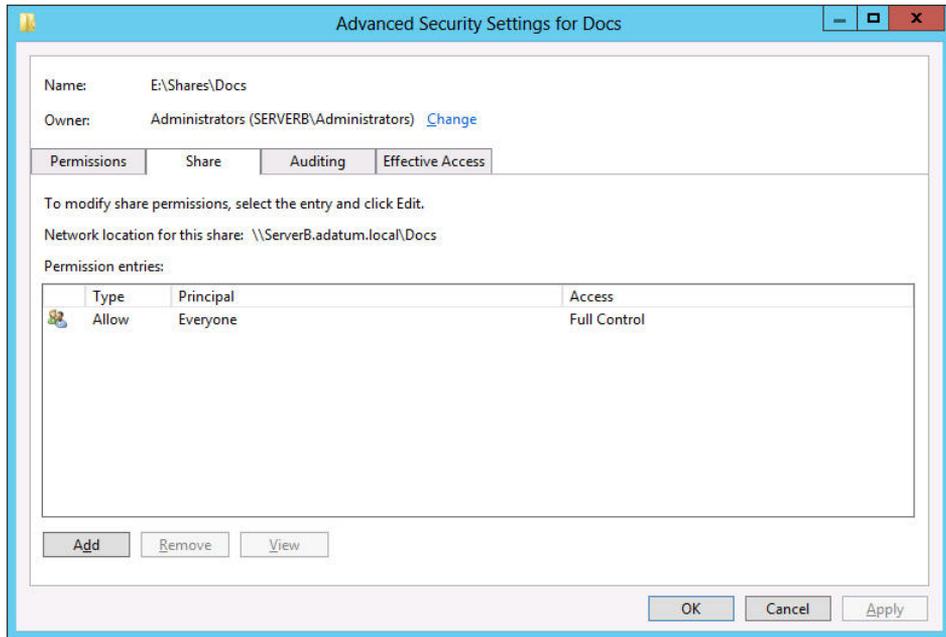


FIGURE 2-8 The Share tab of the Advanced Security Settings dialog box for a share in Server Manager.

7. Click Add to open a Permission Entry dialog box for the share.
8. Click the Select A Principal link to display the Select User, Computer, Service Account, Or Group dialog box.
9. Type the name of or search for the security principal to which you want to assign share permissions and click OK. The security principal you specified appears in the Permission Entry dialog box.
10. Select the type of permissions you want to assign (Allow or Deny).
11. Select the check boxes for the permissions you want to assign and click OK.
12. The new ACE you just created appears in the Advanced Security Settings dialog box.

NOTE BYPASSING SHARE PERMISSIONS

As discussed later in this lesson, many file server administrators simply leave the Allow Full Control share permission to the Everyone special identity in place, essentially bypassing the share permission system, and rely solely on NTFS permissions for granular file system protection.

- 13.** Click OK to close the Advanced Security Settings dialog box.
- 14.** Click OK to close the share's Properties sheet.
- 15.** Close the Server Manager window.

Understanding NTFS authorization

The majority of Windows installations today use the NTFS and ReFS file systems as opposed to FAT32. One of the main advantages of NTFS and ReFS is that they support permissions, which FAT32 does not. As described earlier in this chapter, every file and folder on an NTFS or ReFS drive has an ACL that consists of ACEs, each of which contains a security principal and the permissions assigned to that principal.

In the NTFS permission system, which ReFS also supports, the security principals involved are users and groups, which Windows refers to by using security identifiers (SIDs). When a user attempts to access an NTFS file or folder, the system reads the user's security access token, which contains the SIDs for the user's account and all the groups to which the user belongs. The system then compares these SIDs to those stored in the file or folder's ACEs to determine what access the user should have. This process is called authorization.

Assigning basic NTFS permissions

Most file server administrators work almost exclusively with basic NTFS permissions because there is no need to work directly with advanced permissions for most common access control tasks.

To assign basic NTFS permissions to a shared folder, the options are essentially the same as with share permissions. You can open the folder's Properties sheet in File Explorer and select the Security tab, or you can open a share's Properties sheet in Server Manager, as in the following procedure.

- 1.** Log on to Windows Server 2012 and launch Server Manager.
- 2.** Open the Shares home page.

NOTE NTFS PERMISSIONS

NTFS permissions are not limited to shared folders. Every file and folder on an NTFS volume has permissions. Although this procedure describes the process of assigning permissions to a shared folder, you can open the Properties sheet for any folder in a File Explorer window, click the Security tab, and work with its NTFS permissions in the same way.

3. Open the Properties sheet for a share and click Permissions to open the Permissions page.

NOTE NEW SHARE WIZARD

The New Share Wizard displays this same Permissions interface on its Specify Permissions to Control Access page. The rest of this procedure applies equally well to that page and its subsequent dialog boxes.

4. Click Customize Permissions to open the Advanced Security Settings dialog box for the share, displaying the Permissions tab, as shown in Figure 2-9. This dialog box is as close as the Windows graphical interface can come to displaying the contents of an ACL.

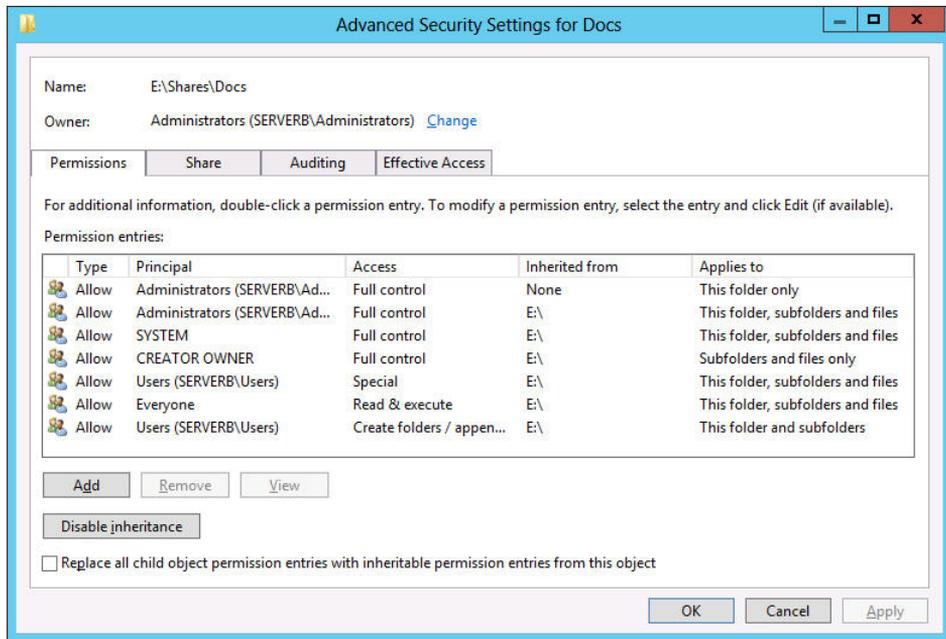


FIGURE 2-9 The Advanced Security Settings dialog box for a share in Server Manager.

5. Click Add. This opens the Permission Entry dialog box for the share.

6. Click the Select A Principal link to display the Select User, Computer, Service Account, or Group dialog box.
7. Type the name of or search for the security principal to which you want to assign share permissions and click OK. The security principal you specified appears in the Permission Entry dialog box.
8. In the Type drop-down list, select the type of permissions you want to assign (Allow or Deny).
9. In the Applies To drop-down list, specify which subfolders and files should inherit the permissions you are assigning.
10. Select the check boxes for the basic permissions you want to assign and click OK. The new ACE you just created appears in the Advanced Security Settings dialog box.
11. Click OK twice to close the Advanced Security Settings dialog box and the Properties sheet.
12. Close the Server Manager window.

Assigning advanced NTFS permissions

In Windows Server 2012, the ability to manage advanced permissions is integrated into the interface you use to manage basic permissions.

In the Permission Entry dialog box, clicking the Show Advanced Permissions link changes the list of basic permissions to a list of advanced permissions. You can then assign advanced permissions in any combination, just as you would basic permissions.

Combining share and NTFS permissions

It is important for file server administrators to understand that the NTFS and share permission systems are completely separate, and that for network users to access files on a shared NTFS drive, they must have both the correct NTFS and the correct share permissions.

The share and NTFS permissions assigned to a file or folder can conflict. For example, if a user has the NTFS Write and Modify permissions for a folder but lacks the share Change permission, that user will not be able to modify a file in that folder.

The share permission system is the simplest of the Windows permission systems, and it provides only basic protection for shared network resources. Share permissions provide only three levels of access, in contrast to the far more complex system of NTFS permissions. Generally, network administrators prefer to use either NTFS or share permissions, not both.

Share permissions provide limited protection, but this might be sufficient on some small networks. Share permissions might also be the only option on a computer with FAT32 drives because the FAT file system does not have its own permission system.

On networks already possessing a well-planned system of NTFS permissions, share permissions are not really necessary. In this case, you can safely leave the Full Control share permission to Everyone, overriding the default Read permission, and allow the NTFS permissions to

provide security. Adding share permissions would complicate the administration process without providing any additional security.

Configuring Volume Shadow Copies

Volume Shadow Copies is a Windows Server 2012 feature that enables you to maintain previous versions of files on a server, so if users accidentally delete or overwrite a file, they can access a copy. You can only implement Shadow Copies for an entire volume; you cannot select specific shares, folders, or files.

To configure a Windows Server 2012 volume to create Shadow Copies, use the following procedure.

1. Log on to Windows Server 2012 using an account with Administrative privileges.
2. Open File Explorer. The File Explorer window appears.
3. In the Folders list, expand the Computer container, right-click a volume and, from the shortcut menu, select Configure Shadow Copies. The Shadow Copies dialog box appears, as shown in Figure 2-10.

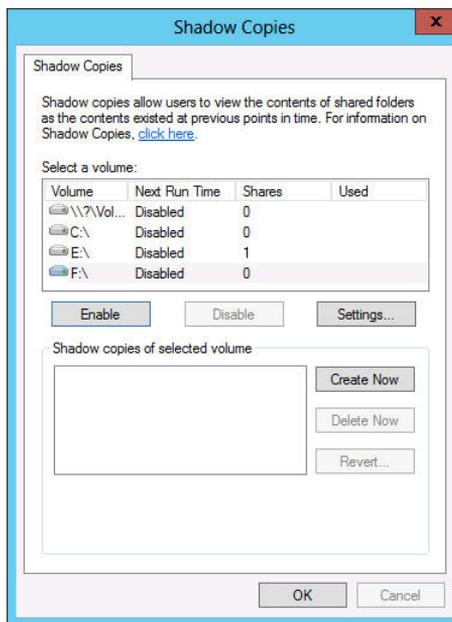


FIGURE 2-10 The Shadow Copies dialog box.

4. In the Select A Volume box, choose the volume for which you want to enable Shadow Copies. By default, when you enable Shadow Copies for a volume, the system uses the following settings:
 - The system stores the shadow copies on the selected volume.

- The system reserves a minimum of 300 MB of disk space for the shadow copies.
 - The system creates shadow copies at 7:00 A.M. and 12:00 P.M. every weekday.
5. To modify the default parameters, click Settings to open the Settings dialog box.
 6. In the Storage Area box, specify the volume where you want to store the shadow copies.
 7. Specify the Maximum Size for the storage area or choose the No Limit option. If the storage area becomes filled, the system begins deleting the oldest shadow copies.
 8. Click Schedule to open the Schedule dialog box. By using the controls provided, you can modify the existing Shadow Copies tasks, delete them, or create new ones, based on the needs of your users.
 9. Click OK twice to close the Schedule and Settings dialog boxes.
 10. Click Enable. The system enables the Shadow Copies feature for the selected volume and creates the first copy in the designated storage area.
 11. Close File Explorer.

After you complete this procedure, users can restore previous versions of files on the selected volumes from the Previous Versions tab on any file or folder's Properties sheet.

Configuring NTFS quotas

Managing disk space is a constant concern for server administrators, and one way to prevent users from monopolizing storage is to implement quotas. Windows Server 2012 supports two types of storage quotas. The more elaborate of the two is implemented as part of File Server Resource Manager. The second, simpler option is NTFS quotas.

NTFS quotas enable administrators to set a storage limit for users of a particular volume. Depending on how you configure the quota, users exceeding the limit can either be denied disk space or just receive a warning. The space consumed by individual users is measured by the size of the files they own or create.

NTFS quotas are relatively limited in that you can set only a single limit for all the users of a volume. The feature is also limited in the actions it can take in response to a user exceeding the limit. The quotas in File Server Resource Manager, by contrast, are much more flexible in the limits you can set and the responses of the program, which can send email notifications, execute commands, and generate reports, as well as log events.

To configure NTFS quotas for a volume, use the following procedure.

1. Log on to Windows Server 2012 using an account with Administrative privileges.
2. Open File Explorer. The File Explorer window appears.
3. In the Folders list, expand the Computer container, right-click a volume and, from the shortcut menu, select Properties. The Properties sheet for the volume appears.
4. Click the Quota tab to display the interface shown in Figure 2-11.

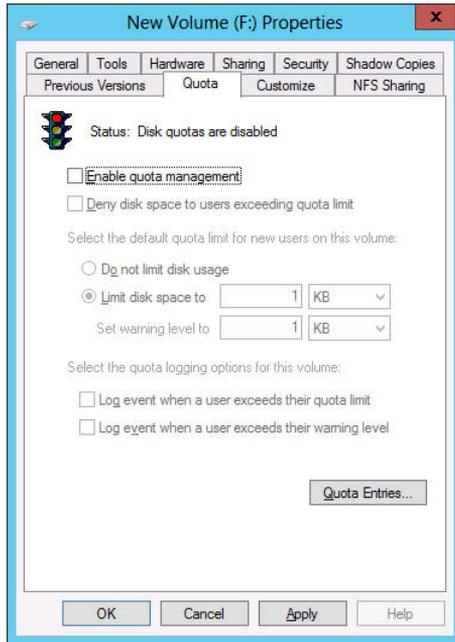


FIGURE 2-11 The Quota tab of a volume's Properties sheet.

5. Select the Enable Quota Management check box to activate the rest of the controls.
6. If you want to prevent users from consuming more than their quota of disk space, select the Deny Disk Space To Users Exceeding Quota Limit check box.
7. Select the Limit Disk Space To option and specify amounts for the quota limit and the warning level.
8. Select the Log Event check boxes to control whether users exceeding the specified limits should trigger log entries.
9. Click OK to create the quota and close the Properties sheet.
10. Close File Explorer.

Objective summary

- Creating folder shares makes the data stored on a file server's disks accessible to network users.
- NTFS permissions enable you to control access to files and folders by specifying the tasks individual users can perform on them. Share permissions provide rudimentary access control for all the files on a network share. Network users must have the proper share and NTFS permissions to access file server shares.

- ABE applies filters to shared folders based on an individual user's permissions to the files and subfolders in the share. Simply put, users who cannot access a particular shared resource are unable to see that resource on the network.
- Offline Files is a Windows feature that enables client systems to maintain local copies of files they access from server shares.
- Volume Shadow Copies is a Windows Server 2012 feature that enables you to maintain previous versions of files on a server, so if users accidentally delete or overwrite a file, they can access a copy.
- NTFS quotas enable administrators to set a storage limit for users of a particular volume.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of this chapter.

1. What is the maximum number of shadow copies a Windows Server 2012 system can maintain for each volume?
 - A. 8
 - B. 16
 - C. 64
 - D. 128
2. Which of the following terms describes the process of granting users access to file server shares by reading their permissions?
 - A. Authentication
 - B. Authorization
 - C. Enumeration
 - D. Assignment
3. Which of the following are tasks you can perform by using the quotas in File Server Resource Manager but can't perform by using NTFS quotas? (Choose all that apply.)
 - A. Send an email message to an administrator when users exceed their limits.
 - B. Specify different storage limits for each user.
 - C. Prevent users from consuming storage space on a volume beyond their allotted limit.
 - D. Generate warnings to users when they approach their allotted storage limit.

4. In the NTFS permission system, combinations of advanced permissions are also known as _____ permissions. (Choose all that apply.)
- A. Special
 - B. Basic
 - C. Share
 - D. Standard
5. Which of the following best defines the role of the security principal in file system permission assignments?
- A. The only person who can access a file that has no permissions assigned to it
 - B. The person responsible for creating permission policies
 - C. The person assigning the permissions
 - D. The person to whom the permissions are assigned



Thought experiment

In the following thought experiment, apply what you've learned about the objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

You are working the help desk for a corporate network and you receive a call from a user named Leo, who is requesting access to the files for a new classified project called Contoso. The Contoso files are stored in a shared folder on a file server, which is locked in a secured underground data storage facility. After verifying that the user has the appropriate security clearance for the project, you create a new group on the file server called `CONTOSO_USERS` and add Leo's user account to that group. Then, you add the `CONTOSO_USERS` group to the access control list for the Trinity folder on the file server and assign the group the following NTFS permissions:

- Allow Modify
- Allow Read & Execute
- Allow List Folder Contents
- Allow Read
- Allow Write

Later, Leo calls you back to tell you that although he is able to access the Contoso folder and read the files stored there, he has been unable to save changes back to the server.

What is the most likely cause of the problem?

Objective 2.2: Configure print and document services

Like the file-sharing functions discussed in the previous section, print device sharing is one of the most basic applications for which local area networks were designed.

This objective covers how to:

- Configure the Easy Print print driver
- Configure Enterprise Print Management
- Configure drivers
- Configure printer pooling
- Configure print priorities
- Configure printer permissions

Deploying a print server

Installing, sharing, monitoring, and managing a single network print device is relatively simple, but when you are responsible for dozens or even hundreds of print devices on a large enterprise network, these tasks can be overwhelming.

Understanding the Windows print architecture

It is important to understand the terms Microsoft uses when referring to the components of the network printing architecture. Printing in Microsoft Windows typically involves the following four components:

- **Print device** A print device is the actual hardware that produces hard-copy documents on paper or other print media. Windows Server 2012 supports both local print devices, which are attached directly to computer ports, and network interface print devices, which are connected to the network either directly or through another computer.
- **Printer** In Windows, a printer is the software interface through which a computer communicates with a print device. Windows Server 2012 supports numerous physical interfaces, including Universal Serial Bus (USB), IEEE 1394 (FireWire), parallel (LPT), serial (COM), Infrared Data Access (IrDA), Bluetooth ports, and network printing services such as lpr, Internet Printing Protocol (IPP), and standard TCP/IP ports.
- **Print server** A print server is a computer (or standalone device) that receives print jobs from clients and sends them to print devices that are either attached locally or connected to the network.
- **Printer driver** A printer driver is a device driver that converts the print jobs generated by applications into an appropriate string of commands for a specific print device.

Printer drivers are designed for a specific print device and provide applications with access to all the print device's features.

NOTE PRINTING NOMENCLATURE

"Printer" and "print device" are the most commonly misused terms in the Windows printing vocabulary. Obviously, many sources use "printer" to refer to the printing hardware. However, in Windows, printer and print device are not equivalent. For example, you can add a printer to a Windows Server 2012 computer without a physical print device being present. The computer can then host the printer, print server, and printer driver. These three components enable the computer to process the print jobs and store them in a print queue until the print device is available.

Understanding Windows printing

These four components work together to process the print jobs produced by Windows applications and turn them into hard-copy documents, as shown in Figure 2-12.

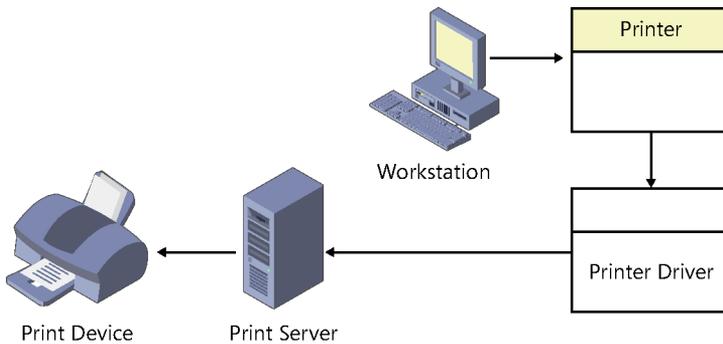


FIGURE 2-12 The Windows print architecture.

Before you can print documents in Windows, you must install at least one printer. To install a printer in Windows, you must do the following:

- Select the print device's specific manufacturer and model.
- Specify the port (or other interface) the computer will use to access the print device.
- Supply a printer driver created specifically for that print device.

When you print a document in an application, you select the printer that will be the destination for the print job.

The printer is associated with a printer driver that takes the commands generated by the application and converts them into a printer control language (PCL), a language understood by the printer. PCLs can be standardized, like the PostScript language, or they can be proprietary languages developed by the print device manufacturer.

The printer driver enables you to configure the print job to use the various capabilities of the print device. These capabilities are typically incorporated into the printer's Properties sheet. For example, your word-processing application does not know if your print device is color or monochrome or if it supports duplex printing. The printer driver provides support for print device features such as these.

After the printer processes a print job, it stores the job in a print queue, known as a spooler. Depending on the arrangement of the printing components, the spooled jobs might be in PCL format, ready to go to the print device, or in an interim format, in which case the printer driver must process the spooled jobs into the PCL format before sending them to the device. If other jobs are waiting to be printed, a new job might wait in the spooler for some time. When the server finally sends the job to the print device, the device reads the PCL commands and produces the hard-copy document.

Windows printing flexibility

The flexibility of the Windows print architecture is manifested in the different ways you can deploy the four printing components. A single computer can perform all the roles (except for the print device, of course), or you can distribute them across the network. The following sections describe four fundamental configurations that are the basis of most Windows printer deployments. You can scale these configurations up to accommodate a network of virtually any size.

DIRECT PRINTING

The simplest print architecture consists of one print device connected to one computer, also known as a locally attached print device, as shown in Figure 2-13. When you connect a print device directly to a Windows Server 2012 computer and print from an application running on that system, the computer supplies the printer, printer driver, and print server functions.

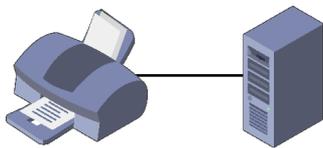


FIGURE 2-13 A locally attached print device.

LOCALLY ATTACHED PRINTER SHARING

In addition to printing from an application running on that computer, you can also share the printer (and the print device) with other users on the same network. In this arrangement, the computer with the locally attached print device functions as a print server. Figure 2-14 shows the other computers on the network, the print clients.

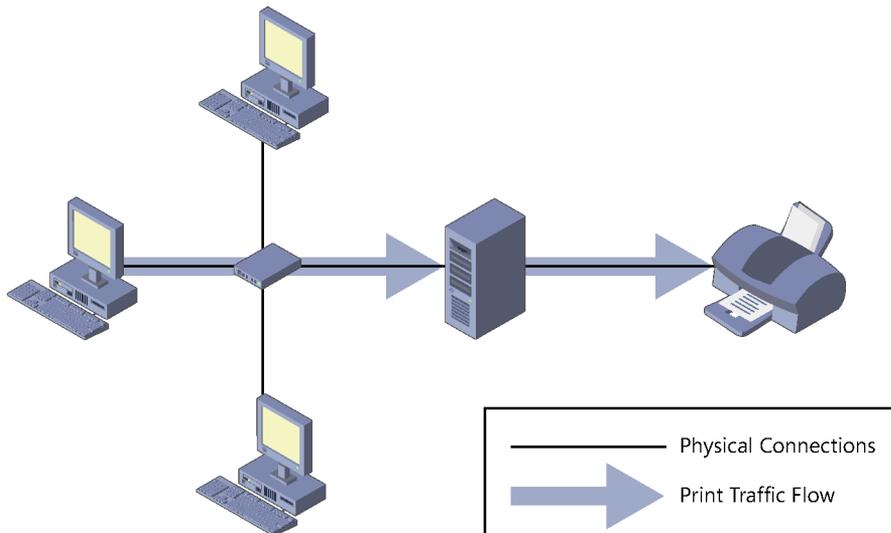


FIGURE 2-14 Sharing a locally attached printer.

In the default Windows Server 2012 printer-sharing configuration, each client uses its own printer and printer driver. As before, the application running on the client computer sends the print job to the printer and the printer driver renders the job, based on the capabilities of the print device.

The main advantage of this printing arrangement is that multiple users, located anywhere on the network, can send jobs to a single print device connected to a computer functioning as a print server. The downside is that processing the print jobs for many users can impose a significant burden on the print server. Although any Windows computer can function as a print server, you should use a workstation for this purpose only when you have no more than a handful of print clients to support or a very light printing volume.

NETWORK-ATTACHED PRINTING

The printing solutions discussed thus far involve print devices connected directly to a computer using a USB or other port. Print devices do not necessarily have to be attached to computers, however. You can connect a print device directly to the network instead. Many print device models are equipped with network interface adapters, enabling you to attach a standard network cable. Some print devices have expansion slots into which you can install a network printing adapter you have purchased separately. Finally, for print devices with no networking capabilities, standalone network print servers are available, which connect to the network and enable you to attach one or more print devices. Print devices so equipped have their own IP addresses and typically have an embedded Web-based configuration interface.

With network-attached print devices, the primary deployment decision the administrator must make is which computer will function as the print server. One simple (but often impractical) option is to let each print client function as its own print server, as shown in Figure 2-15.

Each client processes and spools its own print jobs, connects to the print device by using a TCP (Transmission Control Protocol) port, and sends the jobs directly to the device for printing.

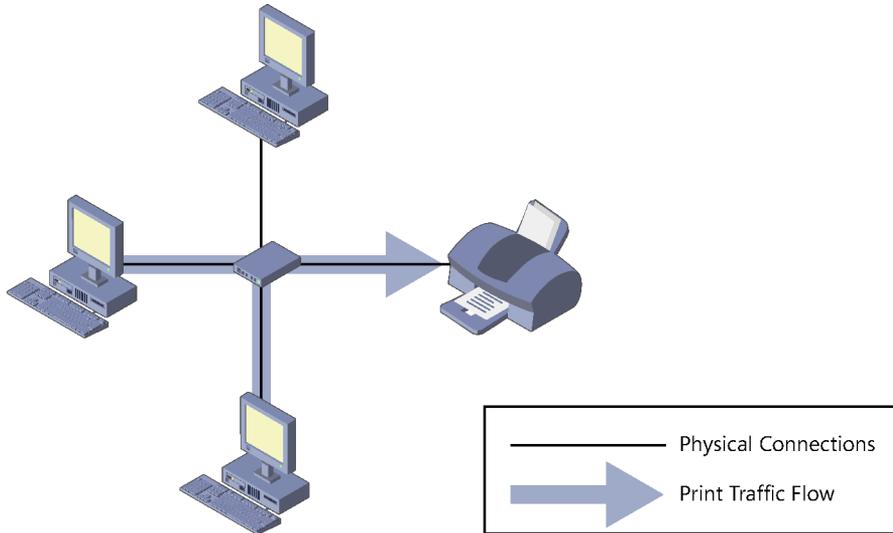


FIGURE 2-15 A network-attached print device with multiple print servers.

Even individual end users with no administrative assistance will find this arrangement simple to set up. However, the disadvantages are many, including the following:

- Users examining the print queue see only their own jobs.
- Users are oblivious of the other users accessing the print device. They have no way of knowing what other jobs have been sent to the print device or how long it will be until the print device completes their jobs.
- Administrators have no way of centrally managing the print queue because each client has its own print queue.
- Administrators cannot implement advanced printing features, such as printer pools or remote administration.
- Error messages appear only on the computer that originated the job the print device is currently processing.
- All print job processing is performed by the client computer rather than being partially offloaded to an external print server.

For these reasons, this arrangement is only suitable for small workgroup networks that do not have dedicated administrators supporting them.

NETWORK-ATTACHED PRINTER SHARING

The other, far more popular option for network-attached printing is to designate one computer as a print server and use it to service all the print clients on the network. To do this, you install a printer on one computer, the print server, and configure it to access the print device directly through a TCP port. You then share the printer, just as you would a locally attached print device, and configure the clients to access the print share.

As you can see in Figure 2-16, the physical configuration is the same as in the previous arrangement, but the logical path the print jobs take on the way to the print device is different. Instead of going straight to the print device, the jobs go to the print server, which spools them and sends them to the print device in order.

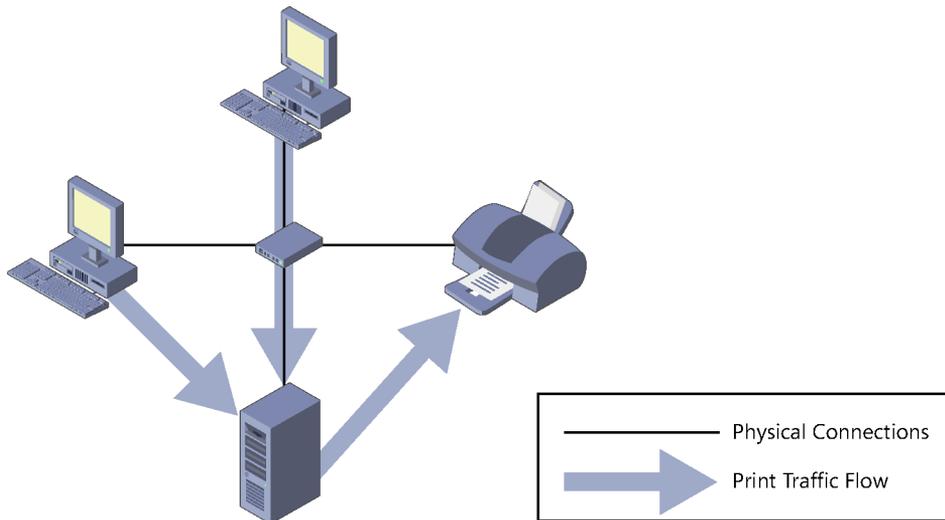


FIGURE 2-16 A network-attached print device with a single shared print server.

With this arrangement, virtually all the disadvantages of the multiple print server arrangement become advantages:

- All the client jobs are stored in a single print queue, so users and administrators can see a complete list of the jobs waiting to be printed.
- Part of the job rendering burden is shifted to the print server, returning control of the client computer to the user more quickly.
- Administrators can manage all the queued jobs from a remote location.
- Print error messages appear on all client computers.
- Administrators can implement printer pools and other advanced printing features.
- Administrators can manage security, auditing, monitoring, and logging functions from a central location.

ADVANCED PRINTING CONFIGURATIONS

Administrators can use the four configurations described in the previous sections as building blocks to create printing solutions for their networks. Many possible variations can be used to create a network printing architecture that supports your organization's needs. Some of the more advanced possibilities are as follows:

- You can connect a single printer to multiple print devices, creating what is called a printer pool. On a busy network with many print clients, the print server can distribute large numbers of incoming jobs among several identical print devices to provide more timely service and better fault tolerance.
- You can connect multiple print devices that support different forms and paper sizes to a single print server, which will distribute jobs with different requirements to the appropriate print devices.
- You can connect multiple print servers to a single print device. By creating multiple print servers, you can configure different priorities, security settings, auditing, and monitoring parameters for different users. For example, you can create a high-priority print server for company executives and a lower-priority print server for junior users. This ensures that the executives' jobs get printed first, even if the servers are connected to the same print device.

Sharing a printer

Using Windows Server 2012 as a print server can be simple or complex, depending on how many clients the server has to support and how much printing they do. For a home or small business network, in which a handful of users need occasional access to the printer, no special preparation is necessary. However, if the computer must support heavy printer use, hardware upgrades, such as additional disk space or system memory, might be needed.

You might also consider making the computer a dedicated print server. In addition to memory and disk space, using Windows Server 2012 as a print server requires processor clock cycles, just like any other application. On a server handling heavy print traffic, other roles and applications are likely to experience substantial performance degradation. If you need a print server to handle heavy traffic, consider dedicating the computer to print server tasks only and deploying other roles and applications elsewhere.

On a Windows Server 2012 computer, you can share a printer as you are installing it or at any time afterward. On older printers, initiate the installation process by launching the Add Printer Wizard from the Printers control panel. However, most of the print devices on the market today use either a USB connection to a computer or an Ethernet connection to a network.

In the case of a USB-connected printer, you plug the print device into a USB port on the computer and turn on the device to initiate the installation process. Manual intervention is required only when Windows Server 2012 does not have a driver for the print device.

For network-attached print devices, an installation program supplied with the product locates the print device on the network, installs the correct drivers, creates a printer on the computer, and configures the printer with the proper IP address and other settings.

After the printer is installed on the Windows Server 2012 computer that will function as your print server, you can share it with your network clients by using the following procedure.

1. Log on to Windows Server.
2. Open the Devices and Printers control panel. The Devices and Printers window appears.
3. Right-click the icon for the printer you want to share and, from the shortcut menu, select Printer Properties. The printer's Properties sheet appears.

NOTE PROPERTIES

The shortcut menu for every printer provides access to two Properties sheets. The Printer Properties menu item opens the Properties sheet for the printer and the Properties menu item opens the Properties sheet for the print device.

4. Click the Sharing tab.
5. Select the Share This Printer check box. The printer name appears in the Share Name text box. You can accept the default name or supply one of your own.
6. Select one or both of the following optional check boxes:
 - **Render Print Jobs On Client Computers** Minimizes the resource utilization on the print server by forcing the print clients to perform the bulk of the print processing.
 - **List In The Directory** Creates a new printer object in the Active Directory Domain Services (AD DS) database, enabling domain users to locate the printer by searching the directory. This option appears only when the computer is a member of an AD DS domain.
7. Click Additional Drivers to open the Additional Drivers dialog box. This dialog box enables you to load printer drivers for other Windows platforms, such as Itanium and x86. When you install the alternate drivers, the print server automatically supplies them to clients running those operating system versions.
8. Select any combination of the available check boxes and click OK. For each check box you select, Windows Server 2012 displays a Printer Drivers dialog box.
9. In each Printer Drivers dialog box, type or browse to the location of the printer drivers for the selected operating system, and then click OK.
10. Click OK to close the Additional Drivers dialog box.
11. Click OK to close the Properties sheet for the printer. The printer icon in the Printers control panel now includes a symbol indicating that it has been shared.

12. Close the control panel.

At this point, the printer is available to clients on the network.

Managing printer drivers

Printer drivers are the components that enable your computers to manage the capabilities of your print devices. When you install a printer on a server running Windows Server 2012, you also install a driver that other Windows computers can use.

Point and Print is the Windows function that enables clients to access the printers installed on print servers. A user on a workstation can select a printer on a server and Windows will automatically install the driver the client needs to process its own print jobs and send them to that printer.

The printer drivers you install on Windows Server 2012 are the same drivers that Windows workstations and other server versions use, with one stipulation. As a 64-bit platform, Windows Server 2012 uses 64-bit device drivers, which are suitable for other computers running 64-bit versions of Windows. If you have 32-bit Windows systems on your network, however, you must install a 32-bit driver on the server for those systems to use.

The Additional Drivers dialog box, accessible from the Sharing tab of a printer's Properties sheet, enables you to install drivers for other processor platforms. However, you must install those drivers from a computer running on the alternative platform. In other words, to install a 32-bit driver for a printer on a server running Windows Server 2012, you must access the printer's Properties sheet from a computer running a 32-bit version of Windows. You can do this by accessing the printer directly through the network by using File Explorer or by running the Print Management snap-in on the 32-bit system and using it to manage your Windows Server 2012 print server.

NOTE INSTALLING DRIVERS

For the server to provide drivers supporting different platforms to client computers, you must make sure when installing the drivers for the same print device that they have the identical name. For example, Windows Server 2012 will treat "HP LaserJet 5200 PCL6" and "HP LaserJet 5200 PCL 6" as two different drivers. The names must be identical for the server to apply them properly.

Using remote access Easy Print

When a Remote Desktop Services client connects to a server, it runs applications using the server's processor(s) and memory. However, if that client wants to print a document from one of those applications, it wants the print job to go to the print device connected to the client computer.

The component that enables Remote Desktop clients to print to their local print devices is called Easy Print. Easy Print takes the form of a printer driver that is installed on the server along with the Remote Desktop Session Host role service.

The Remote Desktop Easy Print driver appears automatically in the Print Management snap-in, but it is not associated with a particular print device. Instead, the driver functions as a redirector, enabling the server to access the printers on the connected clients.

Easy Print requires no configuration other than the installation of the Remote Desktop Services role. However, once it is operational, it provides the server administrator with additional access to the printers on the Remote Desktop clients.

When a Remote Desktop client connects to a server by using the Remote Desktop Connection program or the RD Web Access site, the printers installed on the client system are redirected to the server and appear in the Print Management snap-in as redirected server printers, as shown in Figure 2-17.

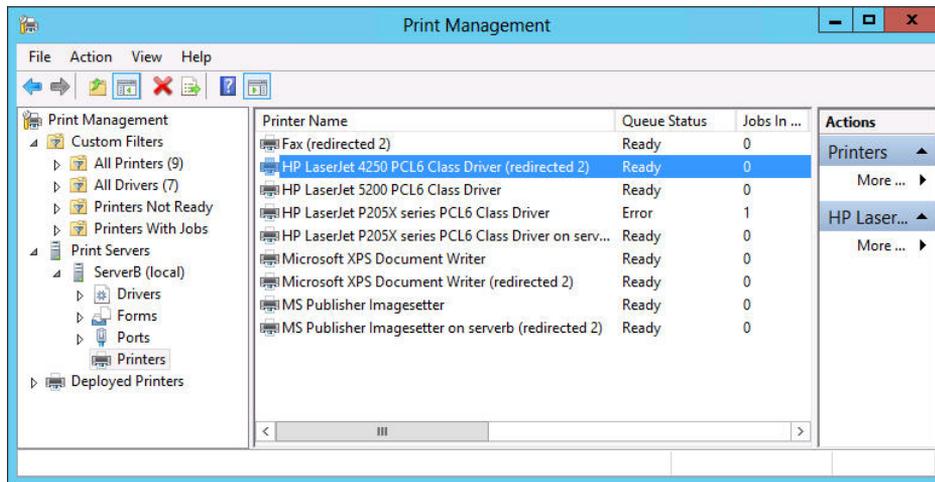


FIGURE 2-17 Printers redirected by Easy Print on a Remote Desktop server.

A client running an application on the server can therefore print to a local print device using the redirected printer. Administrators can also open the Properties sheet for the redirected printer in the usual manner and manipulate its settings.

Configuring printer security

Like folder shares, clients must have the proper permissions to access a shared printer. Printer permissions are much simpler than NTFS permissions; they dictate whether users are allowed to use the printer, manage documents submitted to the printer, or manage the properties of the printer itself. To assign permissions for a printer, use the following procedure.

1. Log on to Windows Server 2012 using a domain account with Administrative privileges.
2. Open Control Panel and select Hardware > Devices and Printers. The Devices and Printers window appears.
3. Right-click one of the printer icons in the window and, from the shortcut menu, select Printer Properties. The printer's Properties sheet appears.

4. Click the Security tab. The top half of the display lists all the security principals currently possessing permissions to the selected printer. The bottom half lists the permissions held by the selected security principal.
5. Click Add. The Select Users, Computers, Or Groups dialog box appears.
6. In the Enter The Object Names To Select text box, type a user or group name, and then click OK. The user or group appears in the Group Or User Names list.
7. Select the security principal you added and select or clear the check boxes in the bottom half of the display to Allow or Deny the user any of the basic permissions.
8. Click OK to close the Properties sheet.
9. Close Control Panel.

Like NTFS permissions, there are two types of printer permissions: basic and advanced. Each of the three basic permissions consists of a combination of advanced permissions.

Managing documents

By default, all printers assign the Allow Print permission to the Everyone special identity, which enables all users to access the printer and manage their own documents. Users who possess the Allow Manage Documents permission can manage any users' documents.

Managing documents refers to pausing, resuming, restarting, and canceling documents that are currently waiting in a print queue. Windows Server 2012 provides a print queue window for every printer, which enables users to view the jobs that are currently waiting to be printed. To manage documents, use the following procedure.

1. Log on to Windows Server 2012.
2. Open Control Panel and select Hardware > Devices and Printers. The Devices and Printers window appears.
3. Right-click one of the printer icons and, from the shortcut menu, select See What's Printing. A print queue window named for the printer appears, as shown in Figure 2-18.

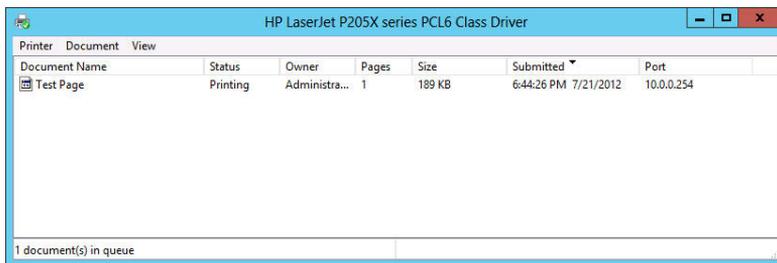


FIGURE 2-18 A Windows Server 2012 print queue window.

4. Select one of the menu items to perform the associated function.
5. Close the print queue window.
6. Close Control Panel.

Managing printers

Users with the Allow Manage This Printer permission can go beyond manipulating queued documents; they can reconfigure the printer itself. Managing a printer refers to altering the operational parameters that affect all users and controlling access to the printer.

Generally, most of the software-based tasks that fall under the category of managing a printer are those you perform once while setting up the printer for the first time. Day-to-day printer management is more likely to involve physical maintenance, such as clearing print jams, reloading paper, and changing toner or ink cartridges. However, the following sections examine some of the printer manager's typical configuration tasks.

Setting printer priorities

In some cases, you might want to give certain users in your organization priority access to a print device so that when print traffic is heavy, their jobs are processed before those of other users. To do this, you must create multiple printers, associate them with the same print device, and then modify their priorities, as described in the following procedure.

1. Log on to Windows Server 2012 using an account with the Manage This Printer permission.
2. Open Control Panel and select Hardware > Devices and Printers. The Devices and Printers window opens.
3. Right-click one of the printer icons and then, from the shortcut menu, select Printer Properties. The Properties sheet for the printer appears.
4. Click the Advanced tab, as shown in Figure 2-19.

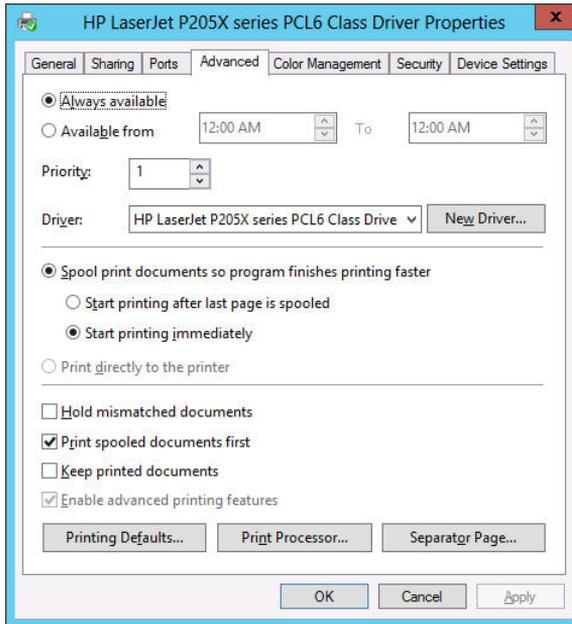


FIGURE 2-19 The Advanced tab of a printer's Properties sheet.

5. Set the Priority spin box to a number representing the highest priority you want to set for the printer. Higher numbers represent higher priorities. The highest possible priority is 99.

NOTE PRINTER PRIORITIES

The values of the Priority spin box do not have any absolute significance; they are pertinent only in relation to one another. As long as one printer has a higher priority value than another, the server will process its print jobs first. In other words, it doesn't matter if the higher priority value is 9 or 99, as long as the lower priority value is less.

6. Click the Security tab.
7. Add the users or groups that you want to provide with high-priority access to the printer and assign the Allow Print permission to them.
8. Revoke the Allow Print permission from the Everyone special identity.
9. Click OK to close the Properties sheet.
10. Create an identical printer using the same printer driver and pointing to the same print device. Leave the Priority setting at its default value of 1 and leave the default permissions in place.
11. Rename the printers, specifying the priority assigned to each one.
12. Close Control Panel.

Inform the privileged users that they should send their jobs to the high-priority printer. All jobs sent to that printer will be processed before those sent to the other, lower-priority printer.

Creating a printer pool

As mentioned earlier, a printer pool increases the production capability of a single printer by connecting it to multiple print devices. When you create a printer pool, the print server sends each incoming job to the first print device it finds that is not busy. This effectively distributes the jobs among the available print devices, providing users with more rapid service.

To configure a printer pool, use the following procedure.

1. Log on to Windows Server 2012 using an account with the Manage Printer permission.
2. Open Control Panel and select Hardware > Devices and Printers. The Devices and Printers window opens.
3. Right-click one of the printer icons and then, from the shortcut menu, select Printer Properties. The Properties sheet for the printer appears.
4. Click the Ports tab.
5. Select all the ports to which the print devices are connected.
6. Select the Enable Printer Pooling check box, and then click OK.
7. Close Control Panel.

To create a printer pool, you must have at least two identical print devices, or at least two print devices that use the same printer driver. The print devices must be in the same location because there is no way to tell which print device will process a given document. You must also connect all the print devices in the pool to the same print server. If the print server is a Windows Server 2012 computer, you can connect the print devices to any viable ports.

Using the Print and Document Services role

All the printer sharing and management capabilities discussed in the previous sections are available on any Windows Server 2012 computer in its default installation configuration. However, installing the Print and Document Services role on the computer provides additional tools that are particularly useful to administrators involved with network printing on an enterprise scale.

When you install the Print and Document Services role by using Server Manager's Add Roles and Features Wizard, a Select Role Services page appears, enabling you to select from the following options:

- **Print Server** Installs the Print Management console for Microsoft Management Console (MMC), which enables administrators to deploy, monitor, and manage printers throughout the enterprise.
- **Distributed Scan Server** Enables the computer to receive documents from network-based scanners and forward them to the appropriate users.

- **Internet Printing** Creates a website that enables users on the Internet to send print jobs to shared Windows printers.
- **LPD Service** Enables UNIX clients running the line printer remote (LPR) program to send their print jobs to Windows printers.

As always, Windows Server 2012 adds a new icon to the Server Manager navigation pane when you install a role. The Print Services home page contains a filtered view of print-related event log entries, a status display for the role-related system services and role services, and performance counters.

The Print Management snap-in for MMC, an administrative tool, consolidates the controls for the printing components throughout the enterprise into a single console. By using this tool, you can access the print queues and Properties sheets for all the network printers in the enterprise, deploy printers to client computers by using Group Policy, and create custom views that simplify the process of detecting print devices that need attention due to errors or depleted consumables.

Windows Server 2012 installs the Print Management console when you add the Print and Document Services role to the computer. You can also install the console without the role by adding the Print and Document Services Tools feature, found under Remote Server Administration Tools > Role Administration Tools in the Add Roles and Features Wizard.

The following sections demonstrate some of the administration tasks you can perform by using the Print Management console.

Adding print servers

By default, the Print Management console displays only the local machine in its list of print servers. Each print server has four nodes beneath it, as shown in Figure 2-20, listing the drivers, forms, ports, and printers associated with that server.

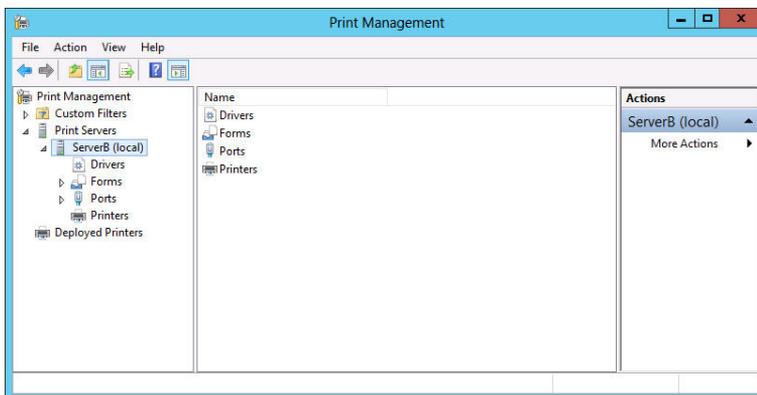


FIGURE 2-20 A print server displayed in the Print Management console.

To manage other print servers and their printers, you must add them to the console by using the following procedure.

1. Log on to Windows Server 2012 and launch Server Manager.
2. Click Tools > Print Management to open the Print Management console.
3. Right-click the Print Servers node and, from the shortcut menu, click Add/Remove Servers to open the Add/Remove Servers dialog box.
4. In the Specify Print Server box, click Browse. The Select Print Server dialog box opens.
5. Select the print server you want to add to the console and click Select Server. The server you selected appears in the Add Server text box in the Add/Remove Servers dialog box.
6. Click Add To List. The server you selected appears in the Print Servers list.
7. Click OK. The server appears under the Print Servers node.
8. Close Control Panel.

You can now manage the printers associated with the server you have added to the console.

Viewing printers

One of the major difficulties for printing administrators on large enterprise networks is keeping track of dozens or hundreds of print devices, all in frequent use and all needing attention on a regular basis. Whether the maintenance required is a major repair, an ink or toner replenishment, or a paper tray refill, print devices will not get the attention they need until an administrator is aware of the problem.

The Print Management console provides multiple ways to view the printing components associated with the print servers on the network. To create views, the console takes the complete list of printers and applies various filters to it, selecting which printers to display. Under the Custom Filters node, there are four default filters, as follows:

- **All Printers** Contains a list of all the printers hosted by all the print servers added to the console
- **All Drivers** Contains a list of all the printer drivers installed on all the print servers added to the console
- **Printers Not Ready** Contains a list of all printers that are not reporting a Ready status
- **Printers With Jobs** Contains a list of all the printers that currently have jobs waiting in the print queue

Views such as Printer Not Ready are a useful way for administrators to identify printers that need attention without having to browse individual print servers or search through a long list of every printer on the network. In addition to these defaults, you can create your own custom filters.

Managing printers and print servers

After you have used filtered views to isolate the printers you want to examine, selecting a printer displays its status, the number of jobs currently in its print queue, and the name of the print server hosting it. If you right-click the filter in the Scope pane and select Show Extended View from the shortcut menu, an additional pane appears containing the contents of the selected printer's queue. You can manipulate the queued jobs just as you would from the Print Queue window in the Print Server console.

The Print Management console also enables administrators to access the configuration interface for any printer or print server appearing in any of its displays. Right-clicking a printer or print server anywhere in the console interface and then selecting Properties from the shortcut menu displays the same Properties sheet you would see on the print server computer itself. Administrators can then configure printers and print servers without having to travel to the site of the print server or establish a Remote Desktop connection to the print server.

Deploying printers with Group Policy

Configuring a print client to access a shared printer is a simple matter of browsing the network or the AD DS tree and selecting the printer. However, when you have to configure hundreds or thousands of print clients, the task becomes more complicated. AD DS helps simplify the process of deploying printers to large numbers of clients.

Publishing printers in the AD DS database enables users and administrators to search for printers by name, location, or model (if you populate the Location and Model fields in the printer object). To create a printer object in the AD DS database, you can either select the List In The Directory check box while sharing the printer or right-click a printer in the Print Management console and, from the shortcut menu, select List In Directory.

To use AD DS to deploy printers to clients, you must configure the appropriate policies in a Group Policy object (GPO). You can link a GPO to any domain, site, or organizational unit (OU) in the AD DS tree. When you configure a GPO to deploy a printer, all the users or computers in that domain, site, or OU will receive the printer connection by default when they log on.

To deploy printers with Group Policy, use the following procedure.

- 1.** Log on to Windows Server 2012 using a domain account with administrative privileges. The Server Manager window opens.
- 2.** Open the Print Management console.
- 3.** Right-click a printer in the console's scope pane and, from the shortcut menu, select Deploy With Group Policy. The Deploy With Group Policy dialog box appears, as shown in Figure 2-21.

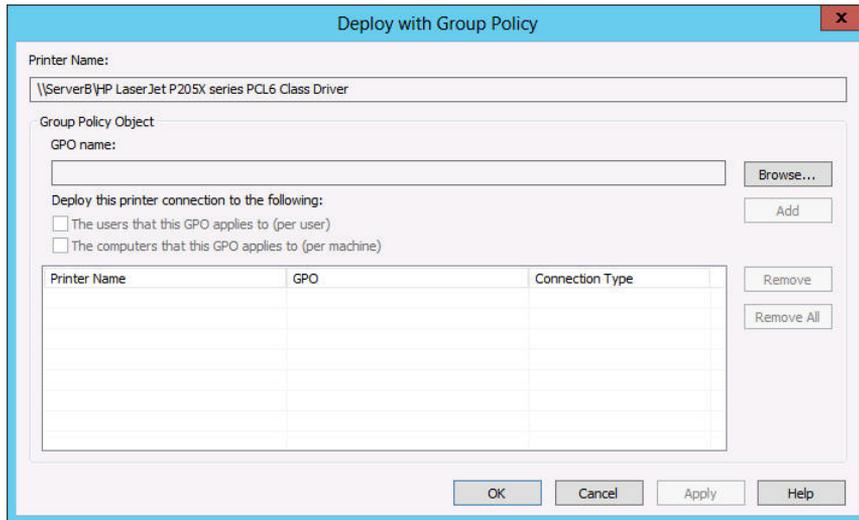


FIGURE 2-21 The Deploy With Group Policy dialog box.

4. Click Browse to open the Browse For A Group Policy Object dialog box.
5. Select the GPO you want to use to deploy the printer and click OK. The GPO you selected appears in the GPO Name field.
6. Select the appropriate check box to select whether to deploy the printer to the users associated with the GPO, the computers, or both, and then click Add. The new printer GPO associations appear in the table.

Deploying the printer to the users means that all the users associated with the GPO will receive the printer connection no matter what computer they use to log on. Deploying the printer to the computers means that all the computers associated with the GPO will receive the printer connection no matter who logs on to them.

7. Click OK. A Print Management message box appears, informing you that the operation has succeeded.
8. Click OK, then click OK again to close the Deploy With Group Policy dialog box.
9. Close Control Panel.

The next time the users running Windows Server 2008 or later and Windows Vista or later who are associated with the GPO refresh their policies or restart, they will receive the new settings and the printer will appear in the Printers control panel.

NOTE PUSHPRINTERCONNECTIONS.EXE

Clients running earlier versions of Windows, including Windows XP and Windows Server 2003, do not support automatic policy-based printer deployments. To enable the GPO to deploy printers on these computers, you must configure the systems to run a utility called PushPrinterConnections.exe. The most convenient way to do this is to configure the same GPO you used for the printer deployment to run the program from a user logon script or machine script.

Objective summary

- Printing in Microsoft Windows typically involves the following four components: print device, printer, print server, and print driver.
- The simplest form of print architecture consists of one print device connected to one computer, known as a locally attached print device. You can share this printer (and the print device) with other users on the same network.
- With network-attached print devices, the administrator's primary deployment decision is which computer will function as the print server.
- Remote Desktop Easy Print is a driver that enables Remote Desktop clients running applications on a server to redirect their print jobs back to their local print devices.
- Printer permissions are much simpler than NTFS permissions; they dictate whether users are allowed to use the printer, manage documents submitted to the printer, or manage the properties of the printer itself.
- The Print Management snap-in for MMC is an administrative tool that consolidates the controls for the printing components throughout the enterprise into a single console.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of this chapter.

1. Which of the following terms describes the software interface through which a computer communicates with a print device?
 - A. Printer
 - B. Print server
 - C. Printer driver
 - D. Print Management snap-in
2. You are setting up a printer pool on a computer running Windows Server 2012. The printer pool contains three identical print devices. You open the Properties dialog box

for the printer and select the Enable Printer Pooling option on the Ports tab. What must you do next?

- A.** Configure the LPT1 port to support three printers.
 - B.** Select or create the ports mapped to the three printers.
 - C.** On the Device Settings tab, configure the installable options to support two additional print devices.
 - D.** On the Advanced tab, configure the priority for each print device so that printing is distributed among the three print devices.
- 3.** One of your print devices is not working properly, and you want to temporarily prevent users from sending jobs to the printer serving that device. What should you do?
- A.** Stop sharing the printer.
 - B.** Remove the printer from Active Directory.
 - C.** Change the printer port.
 - D.** Rename the share.
- 4.** You are administering a computer running Windows Server 2012 configured as a print server. Users in the Marketing group complain that they cannot print documents using a printer on the server. You view the permissions in the printer's properties. The Marketing group is allowed Manage Documents permission. Why can't the users print to the printer?
- A.** The Everyone group must be granted the Manage Documents permission.
 - B.** The Administrators group must be granted the Manage Printers permission.
 - C.** The Marketing group must be granted the Print permission.
 - D.** The Marketing group must be granted the Manage Printers permission.
- 5.** You are administering a print server running Windows Server 2012. You want to perform maintenance on a print device physically connected to the print server. There are several documents in the print queue. You want to prevent the documents from being printed to the printer, but you don't want users to have to resubmit the documents to the printer. What is the best way to do this?
- A.** Open the printer's Properties dialog box, select the Sharing tab, and then select the Do Not Share This Printer option.
 - B.** Open the printer's Properties dialog box and select a port that is not associated with a print device.
 - C.** Open the printer's queue window, select the first document, and then select Pause from the Document window.
 - D.** Open the printer's queue window and select the Pause Printing option from the Printer menu.



Thought experiment

In the following thought experiment, apply what you've learned about the objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

You are a desktop support technician for a law firm with a group of 10 legal secretaries who provide administrative support to the attorneys. All the secretaries use a single, shared, high-speed laser printer that is connected to a dedicated Windows print server. The secretaries print multiple copies of large documents on a regular basis, and although the laser printer is fast, it runs almost constantly. Sometimes the secretaries have to wait 20 minutes or more after submitting a print job for their documents to reach the top of the queue. The office manager has offered to purchase additional printers for the department. However, the secretaries are accustomed to just clicking Print, and don't like the idea of having to examine multiple print queues to determine which has the fewest jobs before submitting a document.

With this in mind, answer the following question:

What can you do to provide the office with a printing solution that will enable the secretaries to utilize additional printers most efficiently?

Objective 2.3: Configure servers for remote management

Windows Server 2012 is designed to facilitate remote server management so administrators rarely if ever have to work directly at the server console. This conserves server resources that can better be devoted to applications.

This objective covers how to:

- Configure WinRM
- Configure down-level server management
- Configure servers for day-to-day management tasks
- Configure multiserver management
- Configure Server Core
- Configure Windows Firewall

Using Server Manager for remote management

Server Manager has been the primary server administration tool for Windows Server ever since Windows Server 2003. The most obvious improvement to the Server Manager tool in Windows Server 2012 is the ability to perform administrative tasks on remote servers and on the local system.

When you log on to a GUI installation of Windows Server 2012 with an administrative account, Server Manager loads automatically, displaying the Welcome tile. The Server Manager interface consists of a navigation pane on the left containing icons representing various views of server resources. Selecting an icon displays a home page in the right pane, which consists of a number of tiles containing information about the resource. The Dashboard page, which appears by default, contains, in addition to the Welcome tile, thumbnails that summarize the other views available in Server Manager. These other views include a page for the Local Server, a page for All Servers, and others for server groups and role groups.

Adding servers

The primary difference between the Windows Server 2012 Server Manager and previous versions is the ability to add and manage multiple servers at once. Although only the local server appears in Server Manager when you first run it, you can add other servers, enabling you to manage them together. The servers you add can be physical or virtual and can be running any version of Windows Server since Windows Server 2003. After you add servers to the interface, you can create groups containing collections of servers, such as those at a particular location or those performing a particular function. These groups appear in the navigation pane, enabling you to administer them as a single entity.

To add servers in Server Manager, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with Administrative privileges. The Server Manager window appears.
2. In the navigation pane, click the All Servers icon to open the All Servers home page.
3. From the Manage menu, select Add Servers to open the Add Servers dialog box.
4. Select one of the following tabs to specify how you want to locate servers to add:
 - **Active Directory** Enables you to search for computers running specific operating systems in specific locations in the local AD DS domain
 - **DNS** Enables you to search for servers in your currently configured Domain Name System (DNS) server
 - **Import** Enables you to supply a text file containing the names or IP addresses of the servers you want to add
5. Initiate a search or upload a text file to display a list of available servers.
6. Select the servers you want to add and click the right arrow button to add them to the Selected list, as shown in Figure 2-22.

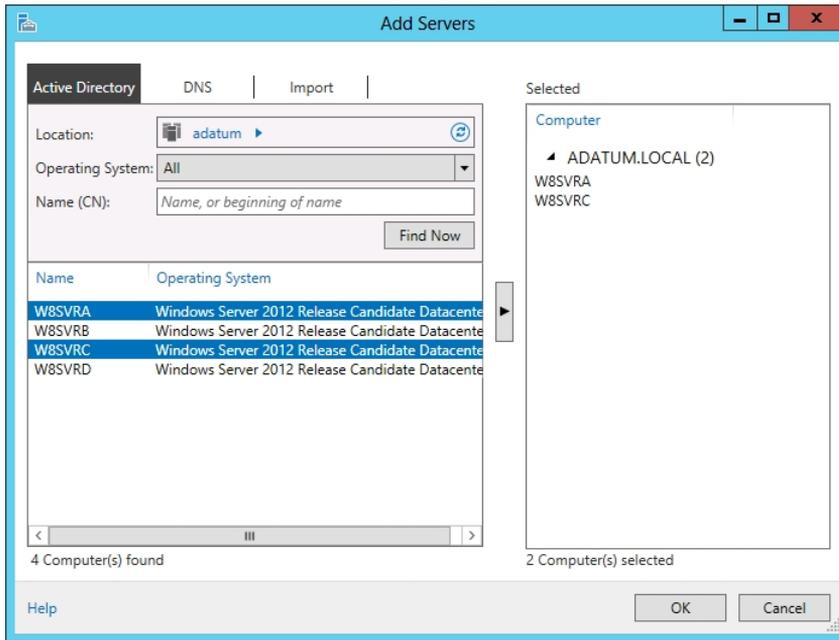


FIGURE 2-22 Selecting servers in Server Manager.

7. Click OK. The servers you selected are added to the All Servers home page.
8. Close the Server Manger console.

Once you have added remote servers to the Server Manager interface, they appear on the All Servers home page. You can then access them in a variety of ways, depending on the version of Windows the remote server is running.

Managing Windows Server 2012 servers

When you add servers running Windows Server 2012 to Server Manager, you can immediately begin using the Add Roles and Features Wizard to install roles and features on any of the servers you have added.

You can also perform other administrative tasks, such as configuring network interface card (NIC) teaming and restarting the server, because Windows Remote Management (WinRM) is enabled by default on Windows Server 2012.

CONFIGURING WINRM

WinRM enables administrators to manage a computer from a remote location by using tools based on Windows Management Instrumentation (WMI) and Windows PowerShell. If the default WinRM setting has been modified, or if you want to change it manually, you can do so through the Server Manager interface.

On the Local Server home page, the Properties tile contains a Remote Management indicator that specifies the server's current WinRM status. To change the WinRM state, click

the Remote Management hyperlink to open the Configure Remote Management dialog box. Clearing the Enable Remote Management Of This Server From Other Computers check box disables WinRM, and selecting the check box enables it.

NOTE USING WINDOWS POWERSHELL

To manage WinRM from a Windows PowerShell session, as in the case of a computer with a Server Core installation, use the following command:

```
Configure-SMRemoting.exe -Get|-Enable|-Disable
```

- **-Get** Displays the current WinRM status
- **-Enable** Enables WinRM
- **-Disable** Disables WinRM

CONFIGURING WINDOWS FIREWALL

However, if you attempt to launch MMC snap-ins targeting a remote server, such as the Computer Management console, you will receive an error because of the default Windows Firewall settings in Windows Server 2012. MMC uses the Distributed Component Object Model (DCOM) for remote management instead of WinRM, and these settings are not enabled by default.

To address this problem, you must enable the following inbound Windows Firewall rules on the remote server you want to manage:

- COM+ Network Access (DCOM-In)
- Remote Event Log Management (NP-In)
- Remote Event Log Management (RPC)
- Remote Event Log Management (RPC-EPMAP)

To modify the firewall rules on the remote system, you can use any one of the following methods:

- Open the Windows Firewall with Advanced Security MMC snap-in on the remote server (if it is a Full GUI installation).
- Run the Netsh AdvFirewall command from an administrative command prompt.
- Use the NetSecurity module in Windows PowerShell.
- Create a GPO containing the appropriate settings and apply it to the remote server.

NOTE USING WINDOWS POWERSHELL

To configure the Windows Firewall rules required for remote server management using DCOM on a Server Core installation, you can use the following Windows PowerShell syntax:

```
Set-NetFirewallRule -name <rule name> -enabled True
```

To obtain the Windows PowerShell names for the preconfigured rules in Windows Firewall, you use the `Get-NetFirewallRule` command. The resulting commands to enable the four rules listed earlier are as follows:

```
Set-NetFirewallRule -name  
  ComPlusNetworkAccess-DCOM-In -enabled True  
Set-NetFirewallRule -name  
  RemoteEventLogSvc-In-TCP -enabled True  
Set-NetFirewallRule -name RemoteEventLogSvc-NP-In-TCP  
-enabled True  
Set-NetFirewallRule -name  
  RemoteEventLogSvc-RPCSS-In-TCP -5  
enabled True
```

For the administrator interested in remote management solutions, the Group Policy method provides distinct advantages. It not only enables you to configure the firewall on the remote system without accessing the server console directly but also can configure the firewall on Server Core installations without having to work from the command line. Finally, and possibly most important for large networks, you can use Group Policy to configure the firewall on all the servers you want to manage at once.

To configure Windows Firewall settings by using Group Policy, use the following procedure. This procedure assumes the server is a member of an AD DS domain and has the Group Policy Management feature installed.

1. Log on to the server running Windows Server 2012 using an account with Administrative privileges. The Server Manager window appears.
2. Open the Group Policy Management console and create a new GPO, giving it a name like Server Firewall Configuration.
3. Open the GPO you created using the Group Policy Management Editor.

MORE INFO GPOS

For more detailed information on creating GPOs and linking them to other objects, see Objective 6.1, "Create Group Policy Objects (GPOs)."

4. Browse to the Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Inbound Rules node.
5. Right-click Inbound Rules and, from the shortcut menu, select New Rule. The New Inbound Rule Wizard appears, displaying the Rule Type page.
6. Select the Predefined option and, in the drop-down list, select COM+ Network Access and click Next. The Predefined Rules page opens.
7. Click Next to open the Action page.
8. Leave the Allow The Connection option selected and click Finish. The rule appears in the Group Policy Management Editor console.

9. Open the New Inbound Rule Wizard again.
10. Select the Predefined option and, in the drop-down list, select Remote Event Log Management. Click Next. The Predefined Rules page opens, displaying the three rules in the Remote Event Log Management group.
11. Leave the three rules selected and click Next to open the Action page.
12. Leave the Allow The Connection option selected and click Finish. The three rules appear in the Group Policy Management Editor console.
13. Close the Group Policy Management Editor console.
14. In the Group Policy Management console, link the Server Firewall Configuration GPO you just created to your domain.
15. Close the Group Policy Management console.

The settings in the GPO you created will be deployed to your remote servers the next time they recycle or restart, and you will be able to use MMC snap-ins, such as Computer Management and Disk Management, on them.

Managing downlevel servers

The Windows Firewall rules you have to enable for remote servers running Windows Server 2012 are also disabled by default on computers running earlier versions of Windows Server, so you also have to enable them there.

Unlike Windows Server 2012, however, earlier versions of the operating system lack the WinRM support needed for them to be managed by using the new Server Manager.

By default, when you add servers running Windows Server 2008 or Windows Server 2008 R2 to the Windows Server 2012 Server Manager, they appear with a manageability status that reads “Online - Verify WinRM 3.0 service is installed, running, and required firewall ports are open.”

To add WinRM support to servers running Windows Server 2008 or Windows Server 2008 R2, you must download and install the following updates:

- .NET Framework 4.0
- Windows Management Framework 3.0

These updates are available from the Microsoft Download Center at the following respective URLs:

- <http://www.microsoft.com/en-us/download/details.aspx?id=17718>
- <http://www.microsoft.com/en-us/download/details.aspx?id=34595>

After you install the updates, the system automatically starts the Windows Remote Management service, but you must still complete the following tasks on the remote server:

- Enable the Windows Remote Management (HTTP-In) rules in Windows Firewall, as shown in Figure 2-23.

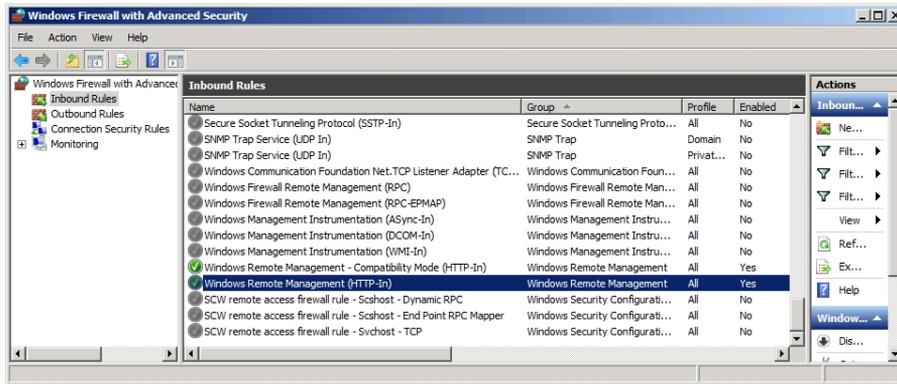


FIGURE 2-23 The Windows Remote Management rules in the Windows Firewall with Advanced Security console.

- Create a WinRM listener by running the `winrm quickconfig` command at a command prompt with Administrative privileges.
- Enable the COM+ Network Access and Remote Event Log Management rules in Windows Firewall, as described in the previous section.

After installing the updates listed here, there are still limitations to the management tasks you can perform on earlier versions of Windows Server from a remote location. For example, you cannot use the Add Roles and Features Wizard in Server Manager to install roles and features on earlier versions of Windows Server. These servers do not appear in the server pool on the Select Destination Server page.

However, you can use Windows PowerShell to install roles and features on servers running Windows Server 2008 and Windows Server 2008 R2 remotely, as in the following procedure.

1. Log on to the server running Windows Server 2012 using an account with Administrative privileges. The Server Manager window appears.
2. Open a Windows PowerShell session with Administrative privileges.
3. Establish a Windows PowerShell session with the remote computer by using the following command:

```
Enter-PSSession <remote server name> -credential <user name>
```

4. Type the password associated with the user name you specified and press Enter.
5. Display a list of the roles and features on the remote server by using the following command:

```
Get-WindowsFeature
```

6. Using the short name of the role or service as it appears in the `Get-WindowsFeature` display, install the component by using the following command.

```
Add-WindowsFeature <feature name>
```

7. Close the session with the remote server by using the following command:

```
Exit-PSession
```

8. Close the Windows PowerShell window.

NOTE WINDOWS POWERSHELL

When you install a role or feature on a remote server by using Windows PowerShell, the installation does not include the role's management tools like a wizard-based installation does. However, you can install the tools along with the role or feature if you include the `IncludeManagementTools` parameter in the `Install-WindowsFeature` command line. Be aware, however, that in the case of a Server Core installation, adding the `IncludeManagementTools` parameter will not install any MMC snap-ins or other graphical tools.

Creating server groups

For administrators of enterprise networks, it might be necessary to add a large number of servers to Server Manager. To avoid having to work with a long scrolling list of servers, you can create server groups based on server locations, functions, or any other organizational paradigm.

When you create a server group, it appears as an icon in the navigation pane, and you can manage the servers in the group just as you would those in the All Servers group.

To create a server group, use the following procedure.

1. Log on to Windows Server 2012 and launch Server Manager.
2. In the navigation pane, click the All Servers icon. The All Servers home page appears.
3. From the Manage menu, select Create Server Group to open the Create Server Group dialog box, as shown in Figure 2-24.

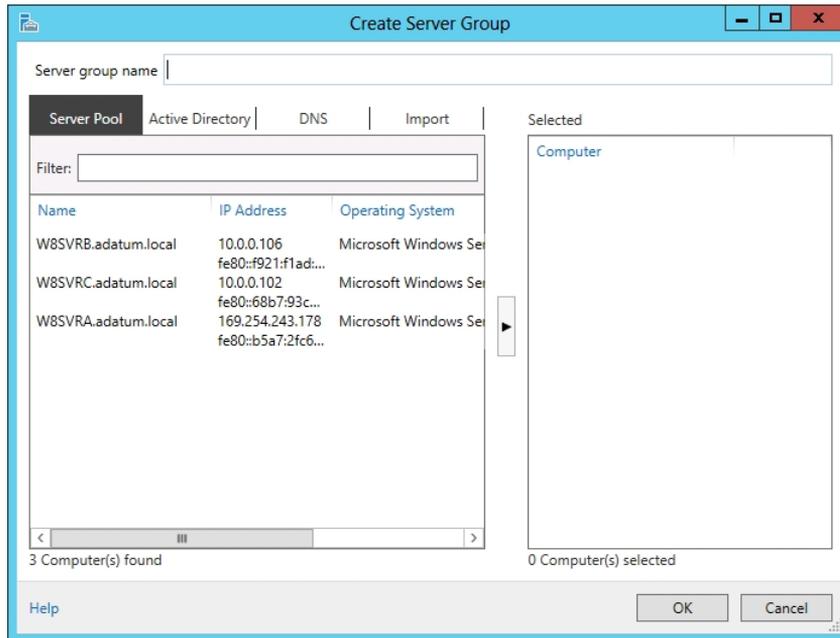


FIGURE 2-24 The Create Server Group dialog box in Server Manager.

4. In the Server Group Name text box, type the name you want to assign to the server group.
5. Select one of the four tabs to choose a method for selecting servers.
6. Select the servers you want to add to the group and click the right arrow button to add them to the Selected box.
7. Click OK. A new server group icon with the name you specified appears in the navigational pane.
8. Close the Server Manager console.

Creating server groups does not affect the functions you can perform on them. You cannot, for example, perform actions on entire groups of servers. The groupings are just a means to keep a large number of servers organized and easy to locate.

Using Remote Server Administration Tools

You can manage remote servers from any computer running Windows Server 2012; all the required tools are installed by default. However, the new administrative method that Microsoft is promoting urges administrators to keep servers locked away and use a workstation to manage servers from a remote location.

To manage Windows servers from a workstation, you must download and install the Remote Server Administration Tools package for the version of Windows running on your workstation from the Microsoft Download Center at <http://www.microsoft.com/download>.

Remote Server Administration Tools is packaged as a Microsoft Update file with an .msu extension, enabling you to deploy it easily from File Explorer, from the command prompt, or by using Software Distribution in a GPO. When you install Remote Server Administration Tools on a workstation running Windows 8, all the tools are activated by default, unlike in previous versions that required you to turn them on by using the Windows Features control panel. You can still use the control panel to turn selected features off, however.

When you launch Server Manager on a Windows workstation, there is no local server and there are no remote servers to manage until you add some. You add servers by using the same process described earlier in this objective.

Your access to the servers you add depends on the account you use to log on to the workstation. If an “Access denied” message appears, you can connect to the server using another account by right-clicking it and, from the shortcut menu, selecting Manage As to display a standard Windows Security dialog box, in which you can supply alternative credentials.

Working with remote servers

Once you have added remote servers to Server Manager, you can access them using a variety of remote administration tools.

Server Manager provides three basic methods for addressing remote servers, as follows:

- **Contextual tasks** When you right-click a server in a Servers tile anywhere in Server Manager, you see a shortcut menu that provides access to tools and commands pointed at the selected server. Some of these are commands that Server Manager executes on the remote server, such as Restart Server and Windows PowerShell. Others launch tools on the local system and direct them at the remote server, such as MMC snap-ins and the Install Roles and Features Wizard. Still others modify Server Manager itself by removing servers from the interface. Other contextual tasks sometimes appear in the Tasks menus for specific panes.
- **Noncontextual tasks** The menu bar at the top of the Server Manager console provides access to internal tasks, such as launching the Add Server and Install Roles and Features Wizards, and the Server Manager Properties dialog box, in which you can specify the console’s refresh interval.
- **Noncontextual tools** The console’s Tools menu provides access to external programs, such as MMC snap-ins and the Windows PowerShell interface, that are directed at the local system.

Objective summary

- Windows Server 2012 is designed to facilitate remote server management so administrators rarely if ever have to work directly at the server console. This conserves server resources that can better be devoted to applications.

- When you add servers running Windows Server 2012 to Server Manager, you can immediately begin using the Add Roles and Features Wizard to install roles and features on any of the servers you have added.
- The Windows Firewall rules you have to enable for remote servers running Windows Server 2012 are also disabled by default on computers running earlier versions of Windows Server, so you also have to enable them there.
- For administrators of enterprise networks, it might be necessary to add a large number of servers to Server Manager. To avoid having to work with a long scrolling list of servers, you can create server groups based on server locations, functions, or any other organizational paradigm.
- You can manage remote servers from any computer running Windows Server 2012; all the required tools are installed by default. However, the new administrative method that Microsoft is promoting urges administrators to keep servers locked away and use a workstation to manage servers from a remote location.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. Which of the following tasks must you perform before you can manage a remote server running Windows Server 2012 using the Computer Management snap-in?
 - A. Enable WinRM on the remote server.
 - B. Enable the COM+ Network Access rule on the remote server.
 - C. Enable the Remote Event Log Management rules on the remote server.
 - D. Install Remote Server Administration Tools on the remote server.
2. Which of the following Windows PowerShell cmdlets can you use to list the existing Windows Firewall rules on a computer running Windows Server 2012? (Choose all that apply.)
 - A. Get-NetFirewallRule
 - B. Set-NetFirewallRule
 - C. Show-NetFirewallRule
 - D. New-NetFirewallRule
3. Which of the following tasks can you NOT perform remotely on a server running Windows Server 2008?
 - A. Install roles by using Server Manager.
 - B. Install roles by using Windows PowerShell.
 - C. Connect to the remote server by using the Computer Management snap-in.
 - D. Monitor event log entries.

4. Which of the following updates must you install on a server running Windows Server 2008 before you can connect to it by using Windows Server 2012 Server Manager? (Choose all that apply.)
 - A. .NET Framework 3.5
 - B. .NET Framework 4.0
 - C. Windows Management Framework 3.0
 - D. Windows Server 2008 R2
5. When you run Server Manager from a Windows 8 workstation using Remote Server Administration Tools, which of the following elements do NOT appear in the default display?
 - A. The Dashboard
 - B. The Local Server home page
 - C. The All Servers home page
 - D. The Welcome tile



Thought experiment

In the following thought experiment, apply what you've learned about the objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

Ralph is responsible for the 24 servers running a particular application, which are scattered across his company's enterprise network. Ralph wants to use Server Manager on his Windows 8 workstation to manage those servers and monitor the events that occur on them. To do this, he must enable the incoming COM+ Network Access and Remote Event Log Management rules in Windows Firewall on the servers.

Because he can't travel to the locations of all the servers and many of the sites do not have trustworthy IT personnel, Ralph has decided to use Group Policy to configure Windows Firewall on all the servers. The company's Active Directory Domain Services tree is organized geographically, which means that Ralph's servers are located in many different OUs under one domain.

With this in mind, answer the following question:

How can Ralph use Group Policy to deploy the required Windows Firewall rule settings to his 24 servers, and only those servers?

Answers

This section contains the solutions to the thought experiments and answers to the lesson review questions in this chapter.

Objective 2.1: Review

1. **Correct answer:** C
 - A. **Incorrect:** Windows Server 2012 can maintain more than 8 volume shadow copies.
 - B. **Incorrect:** Windows Server 2012 can maintain more than 16 volume shadow copies.
 - C. **Correct:** Windows Server 2012 can maintain up to 64 volume shadow copies before it begins deleting the oldest data.
 - D. **Incorrect:** Windows Server 2012 cannot maintain 128 volume shadow copies.
2. **Correct answer:** B
 - A. **Incorrect:** Authentication is the process of verifying the user's identity.
 - B. **Correct:** Authorization is the process by which a user is granted access to specific resources based on the permissions he or she possesses.
 - C. **Incorrect:** Access-based enumeration is a Windows feature that prevents users from seeing resources to which they do not have permissions.
 - D. **Incorrect:** Assignment describes the process of granting permissions, but not reading them.
3. **Correct answers:** A, B
 - A. **Correct:** Using File Server Resource Manager, you can notify administrators with email messages when users exceed their allotment of storage.
 - B. **Correct:** Using File Server Resource Manager, you can create quotas for individual users that specify different storage limits.
 - C. **Incorrect:** You can use NTFS quotas to prevent users from consuming storage space on a volume beyond their allotted limit.
 - D. **Incorrect:** You can use NTFS quotas to generate warnings to users when they approach their allotted storage limit.
4. **Correct answers:** B, D
 - A. **Incorrect:** In Windows Server versions prior to Windows Server 2012, special permissions are combined to form standard permissions.
 - B. **Correct:** Basic permissions are formed by creating various combinations of advanced permissions.

- C. Incorrect:** Share permissions are a system that is separate from the NTFS permission system.
 - D. Correct:** In Windows Server versions prior to Windows Server 2012, standard permissions are formed by creating various combinations of special permissions.
- 5. Correct answer: D**
- A. Incorrect:** The owner is the only person who can access a file that has no permissions assigned to it.
 - B. Incorrect:** The security principal is not the person responsible for creating an organization's permission policies.
 - C. Incorrect:** The security principal receives permissions; the security principal does not create them.
 - D. Correct:** The security principal is the user or computer to which permissions are assigned.

Objective 2.1: Thought experiment

The most likely cause of the problem is that Leo does not have sufficient share permissions for read/write access to the Contoso files. Granting the CONTOSO_USERS group the Allow Full Control share permission should enable Leo to save his changes to the Contoso files.

Objective 2.2: Review

- 1. Correct answer: A**
- A. Correct:** In Windows, a printer is the software interface through which a computer communicates with a print device.
 - B. Incorrect:** A print server is a device that receives print jobs from clients and sends them to print devices that are either attached locally or connected to the network.
 - C. Incorrect:** A printer driver is a device driver that converts the print jobs generated by applications into an appropriate string of commands for a specific print device.
 - D. Incorrect:** The Print Management snap-in is a tool that administrators can use to manage printers all over the network.
- 2. Correct answer: B**
- A. Incorrect:** Whether the printers are pooled or not, each one must be connected to a separate port.
 - B. Correct:** To set up printer pooling, select the Enable Printer Pooling check box, and then select or create the ports corresponding to printers that will be part of the pool.

- C. Incorrect:** You do not use the installable options settings to create a printer pool.
 - D. Incorrect:** Priorities have nothing to do with printer pooling.
- 3. Correct answer: A**
- A. Correct:** If you stop sharing the printer, users will no longer be able to use the print device.
 - B. Incorrect:** Removing the printer from Active Directory will prevent users from finding the printer by using a search, but they can still access it.
 - C. Incorrect:** Changing the printer port will prevent the printer from sending jobs to the print device, but it will not prevent users from sending jobs to the printer.
 - D. Incorrect:** Renaming the share can make it difficult for users to find the printer, but they can still use it when they do find it.
- 4. Correct answer: C**
- A. Incorrect:** The Manage Documents permission does not allow users to send jobs to the printer.
 - B. Incorrect:** The Manage Printers permission does not allow users to send jobs to the printer.
 - C. Correct:** The Print permission allows users to send documents to the printer; the Manage Documents permission does not.
 - D. Incorrect:** The Manage Documents permission does not allow users to send jobs to the printer.
- 5. Correct answer: D**
- A. Incorrect:** A printer that is not shared will continue to process jobs that are already in the queue.
 - B. Incorrect:** Changing the port will require the users to resubmit the jobs that were in the queue.
 - C. Incorrect:** Pausing the first document in the queue will not prevent the other queued jobs from printing.
 - D. Correct:** When you select the Pause Printing option, the documents will remain in the print queue until you resume printing. This option applies to all documents in the queue.

Objective 2.2: Thought experiment

Install additional, identical printers, connecting them to the same Windows Vista print server, and create a printer pool by selecting the appropriate check box on the Ports tab of the printer's Properties sheet.

Objective 2.3: Review

1. Correct answer: B

- A. Incorrect:** WinRM is enabled by default on Windows Server 2012.
- B. Correct:** The COM+ Network Access rule must be enabled on the remote server for MMC snap-ins to connect.
- C. Incorrect:** The Remote Event Log Management rules are not necessary to connect to a remote server using an MMC snap-in.
- D. Incorrect:** PTR records contain the information needed for the server to perform reverse name lookups.

2. Correct answers: A, C

- A. Correct:** The Get-NetFirewallRule cmdlet displays a list of all the rules on a system running Windows Firewall.
- B. Incorrect:** The Set-NetFireWallRule cmdlet is for managing specific rules, not listing them.
- C. Correct:** The Show-NetFirewallRule cmdlet displays a list of all the rules on a system running Windows Firewall.
- D. Incorrect:** The New-NetFireWallRule cmdlet is for creating rules, not listing them.

3. Correct answer: A

- A. Correct:** You cannot install roles on a remote server running Windows Server 2008 by using Server Manager.
- B. Incorrect:** You can install roles on a remote server running Windows Server 2008 by using Windows PowerShell.
- C. Incorrect:** You can connect to a remote server running Windows Server 2008 by using the Computer Management console as long as you enable the COM+ Network Access rule.
- D. Incorrect:** You can monitor event log entries on a remote server running Windows Server 2008 as long as you enable the Remote Event Log Management rules.

4. Correct answers: B, C

- A. Incorrect:** .NET Framework 3.5 is not needed for Server Manager to connect to Windows Server 2008.
- B. Correct:** .NET Framework 4.0 is needed for Server Manager to connect to Windows Server 2008.
- C. Correct:** Windows Management Framework 3.0 is needed for Server Manager to connect to Windows Server 2008.
- D. Incorrect:** It is not necessary to upgrade to Windows Server 2008 R2 for Server Manager to connect to Windows Server 2008.

5. Correct answer: B

- A. Incorrect:** The Dashboard does appear in the default Server Manager display.
- B. Correct:** The Local Server home page does not appear, because the local system is a workstation, not a server.
- C. Incorrect:** The All Servers home page does appear in the default Server Manager display.
- D. Incorrect:** The Welcome tile does appear in the default Server Manager display.

Objective 2.3: Thought experiment

After creating a GPO containing the required Windows Firewall settings, Ralph should create a security group containing all the 24 computer objects representing his servers. Then, he should link the GPO to the company domain and use security filtering to limit the scope of the GPO to the group he created.



Configure Hyper-V

The concept of virtualizing servers has, in the past several years, grown from a novel experiment to a convenient lab and testing tool to a legitimate deployment strategy for production servers. Windows Server 2012 includes the Hyper-V role, which enables administrators to create virtual machines (VMs), each of which runs in its own isolated environment. VMs are self-contained units that administrators can easily move from one physical computer to another, greatly simplifying the process of deploying network applications and services.

This chapter covers some of the fundamental tasks that administrators perform to create and deploy Hyper-V servers and VMs.

Objectives in this chapter:

- Objective 3.1: Create and configure virtual machine settings
- Objective 3.2: Create and configure virtual machine storage
- Objective 3.3: Create and configure virtual networks

Objective 3.1: Create and configure virtual machine settings

Server virtualization in Windows Server 2012 is based on a module called a hypervisor. Sometimes called a virtual machine monitor (VMM), the hypervisor is responsible for abstracting the computer's physical hardware and creating multiple virtualized hardware environments, called VMs. Each VM has its own (virtual) hardware configuration and can run a separate copy of an operating system (OS). Therefore, with sufficient physical hardware and the correct licensing, a single computer running Windows Server 2012 with the Hyper-V role installed can support multiple VMs, which administrators can manage as if they were standalone computers.

This objective covers how to:

- Configure dynamic memory
- Configure smart paging
- Configure Resource Metering
- Configure guest integration services

Virtualization architectures

Virtualization products can use several different architectures to share a computer's hardware resources among VMs. The earlier type of virtualization products, including Microsoft Windows Virtual PC and Microsoft Virtual Server, requires a standard OS installed on a computer. This becomes the "host" OS. Then you install the virtualization product, which adds the hypervisor component. The hypervisor essentially runs alongside the host OS, as shown in Figure 3-1, and enables you to create as many VMs as the computer has hardware to support.

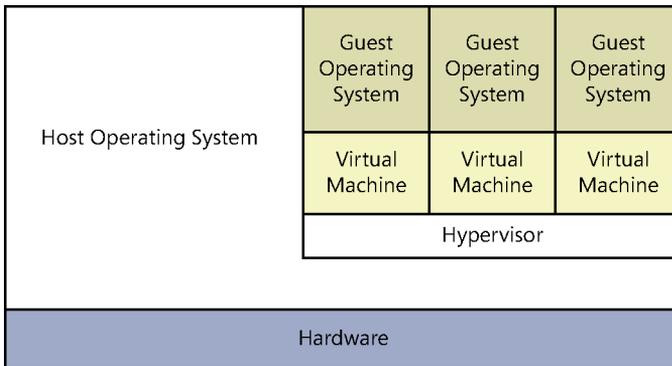


FIGURE 3-1 A hybrid VMM sharing hardware access with a host operating system.

This arrangement, in which the hypervisor runs on top of a host OS, is called Type II virtualization. By using the Type II hypervisor, you create a virtual hardware environment for each VM. You can specify how much memory to allocate to each VM, create virtual disk drives by using space on the computer's physical drives, and provide access to peripheral devices. You then install a "guest" OS on each VM, just as if you were deploying a new computer. The host OS then shares access to the computer's processor with the hypervisor, with each taking the clock cycles it needs and passing control of the processor back to the other.

Type II virtualization can provide adequate VM performance, particularly in classroom and laboratory environments, but it does not provide performance equivalent to separate physical computers. Therefore, it is not generally recommended for high-traffic servers in production environments.

The virtualization capability built into Windows Server 2012, called Hyper-V, uses a different type of architecture. Hyper-V uses Type I virtualization, in which the hypervisor is an abstraction layer that interacts directly with the computer's physical hardware—that is, without an intervening host OS. The term *hypervisor* is intended to represent the level beyond the term *supervisor*, in regard to responsibility for allocating a computer's processor clock cycles.

The hypervisor creates individual environments called partitions, each of which has its own OS installed and accesses the computer's hardware via the hypervisor. Unlike Type II virtualization, no host OS shares processor time with the hypervisor. Instead, the hypervisor designates the first partition it creates as the parent partition and all subsequent partitions as child partitions, as shown in Figure 3-2.

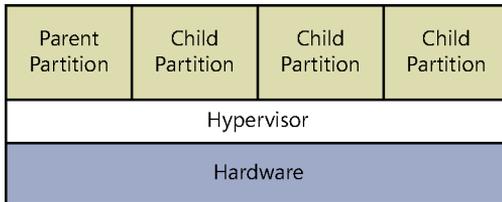


FIGURE 3-2 A Type I VMM, with the hypervisor providing all hardware access.

The parent partition accesses the system hardware through the hypervisor, just as the child partitions do. The only difference is that the parent runs the virtualization stack, which creates and manages the child partitions. The parent partition is also responsible for the subsystems that directly affect the performance of the computer's physical hardware, such as Plug and Play, power management, and error handling. These subsystems also run in the OSs on the child partitions, but they address only virtual hardware, whereas the parent, or root, partition handles the actual hardware.

NOTE HYPER-V

It might not seem like the Hyper-V role in Windows Server 2012 provides Type I virtualization, because it requires the Windows Server OS to be installed and running. However, adding the Hyper-V role actually converts the installed instance of Windows Server 2012 into the parent partition and causes the system to load the hypervisor before the OS.

Hyper-V implementations

Windows Server 2012 includes the Hyper-V role only in the Standard and Datacenter editions. The Hyper-V role is required for the OS to function as a computer's primary partition, enabling it to host other VMs. No special software is required for an OS to function as a guest OS in a VM. Therefore, although Windows Server 2012 Essentials does not include the Hyper-V role, it can function as a guest OS. Other guest OSs supported by Hyper-V include the current Windows workstation OSs and many other non-Microsoft server and workstation products.

Hyper-V licensing

As far as Hyper-V is concerned, the primary difference between the Standard and Datacenter editions of Windows Server 2012 is the number of VMs they support. When you install a Windows Server 2012 instance on a VM, you must have a license for it, just like when you install it on a physical machine. Purchasing the Datacenter edition licenses you to create an unlimited number of VMs running Windows Server 2012 on that one physical machine. The Standard license provides only two virtual instances of Windows Server 2012.

NOTE LICENSING

The licensing restrictions of the Windows Server 2012 Standard and Datacenter editions do not govern how many VMs you can create. They only govern what OS you are permitted to install on the VMs. You can, for example, use a Standard edition license to create only two virtual instances of Windows Server 2012, but you can also create any number of VMs running a free Linux distribution.

Hyper-V hardware limitations

The Windows Server 2012 version of Hyper-V contains massive improvements in the scalability of the system over previous versions. A Windows Server 2012 Hyper-V host system can have up to 320 logical processors, supporting up to 2,048 virtual CPUs and up to 4 terabytes (TB) of physical memory.

One server can host as many as 1,024 active VMs, and each VM can have up to 64 virtual CPUs and up to 1 TB of memory.

Hyper-V can also support clusters with up to 64 nodes and 8,000 VMs.

NOTE WINDOWS POWERSHELL

Another major improvement in the Windows Server 2012 version of Hyper-V is the inclusion of a Hyper-V module for Windows PowerShell, which includes more than 160 new cmdlets dedicated to the creation and management of the Hyper-V service and its VMs.

Hyper-V Server

In addition to the Hyper-V implementation in Windows Server 2012, Microsoft provides a dedicated Hyper-V Server product, which is a subset of Windows Server 2012. Hyper-V Server includes the Hyper-V role, which it installs by default during the OS installation. With the exception of some limited File and Storage Services and Remote Desktop capabilities, the OS includes no other roles, as shown in Figure 3-3.

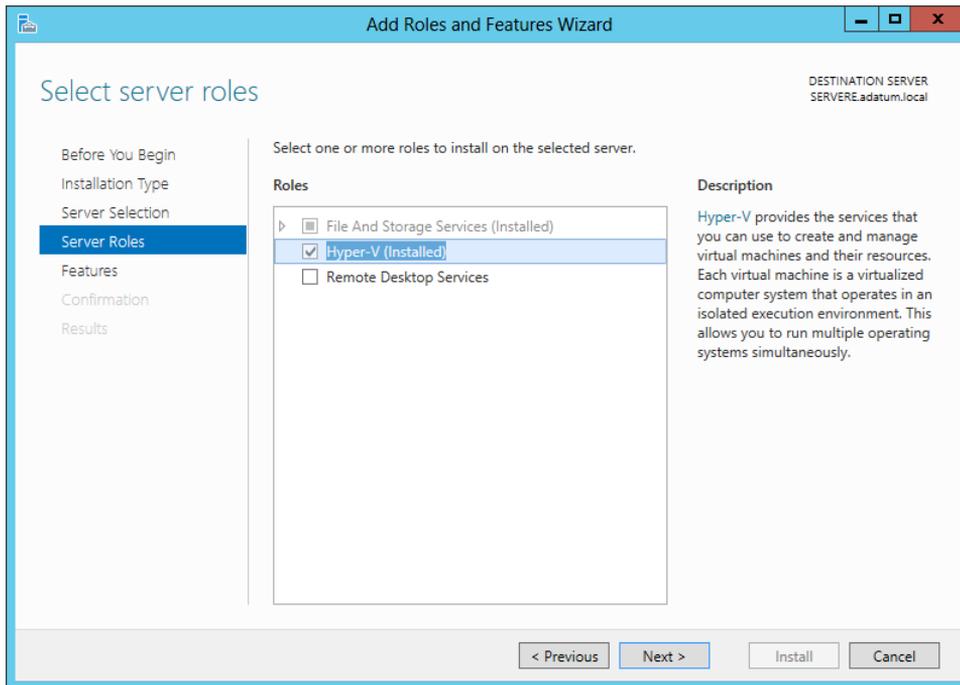


FIGURE 3-3 Roles available in Hyper-V Server.

The Hyper-V Server is also limited to the Server Core interface, although it includes a simple, script-based configuration interface, as shown in Figure 3-4. You can manage Hyper-V Server remotely by using Server Manager and Hyper-V Manager, just as you would any other Server Core installation.

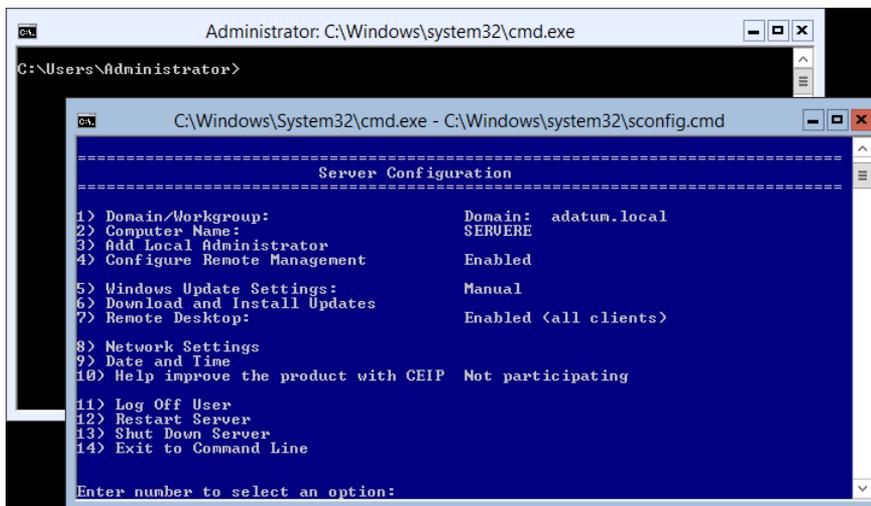


FIGURE 3-4 The Server Core interface in Hyper-V Server.

Unlike Windows Server 2012, Hyper-V Server is a free product, available for download from Microsoft's website. However, Hyper-V Server does not include any licenses for virtual instances. You must obtain and license all the OSs you install on the VMs you create.

Installing Hyper-V

Once you have the appropriate hardware and the required licenses, you can add the Hyper-V role to Windows Server 2012 by using Server Manager, just as you would any other role.

Adding the Hyper-V role installs the hypervisor software, and, in the case of a full GUI installation, also installs the management tools. The primary tool for creating and managing VMs and their components on Hyper-V servers is the Hyper-V Manager console. Hyper-V Manager provides administrators with a list of all the VMs on Windows Server 2012 systems and enables administrators to configure the environments of both the servers and the individual VMs. There is also a set of Hyper-V cmdlets for Windows PowerShell that enables you to exercise complete control over VMs using that interface.

Microsoft recommends that you do not install other roles with Hyper-V. It is better to implement any other roles that you need the physical computer to perform within one of the VMs you create by using Hyper-V. In addition, you might want to consider installing Hyper-V on a computer by using the Server Core installation option. This will minimize the overhead expended on the partition. As with other roles, installing Hyper-V on Server Core excludes the management tools, which you must install separately as a feature on another computer.

Before you can install the Hyper-V role on a server running Windows Server 2012, you must have appropriate hardware, as follows:

- 64-bit processor that includes hardware-assisted virtualization. This is available in processors that include a virtualization option, such as Intel Virtualization Technology (Intel VT) and AMD Virtualization (AMD-V) technology.
- A system BIOS that supports the virtualization hardware, on which the virtualization feature has been enabled.
- Hardware-enforced Data Execution Prevention (DEP), which Intel describes as eXecuted Disable (XD) and AMD describes as No eXecute (NS). This is a technology used in CPUs to segregate areas of memory for either storage of processor instructions or for storage. Specifically, you must enable the Intel XD bit (execute disable bit) or the AMD NX bit (no execute bit).

To install the Hyper-V role, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Manage menu, select Add Roles And Features. The Add Roles and Features Wizard starts, displaying the Before You Begin page.
3. Click Next to open the Select Installation Type page.

4. Leave the Role-Based Or Feature-Based Installation option selected and click Next. The Select Destination Server page opens.
5. Select the server on which you want to install Hyper-V and click Next. The Select Server Roles page opens.
6. Select the Hyper-V role. The Add Features That Are Required for Hyper-V dialog box appears.
7. Click Add Features to accept the dependencies, and then click Next to open the Select Features page.
8. Click Next to open the Hyper-V page.
9. Click Next. The Create Virtual Switches page opens, as shown in Figure 3-5.

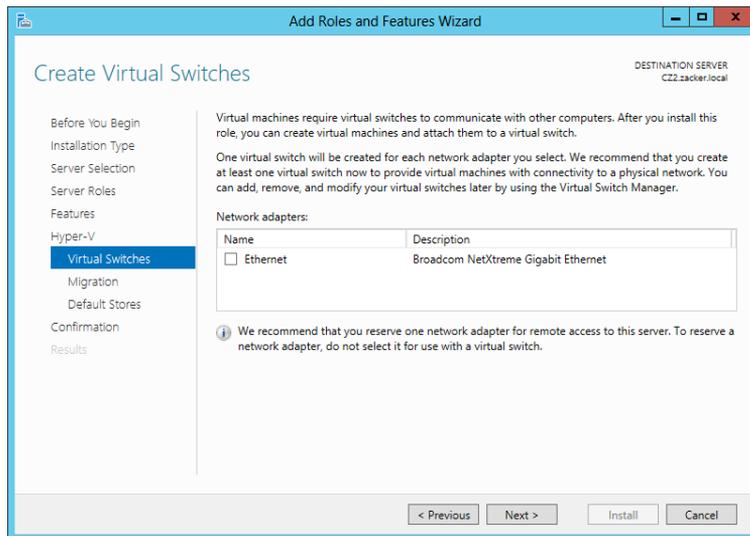


FIGURE 3-5 The Create Virtual Switches page of the Add Roles and Features Wizard.

10. Select the appropriate check box for a network adapter and click Next. The Virtual Machine Migration page opens, as shown in Figure 3-6.

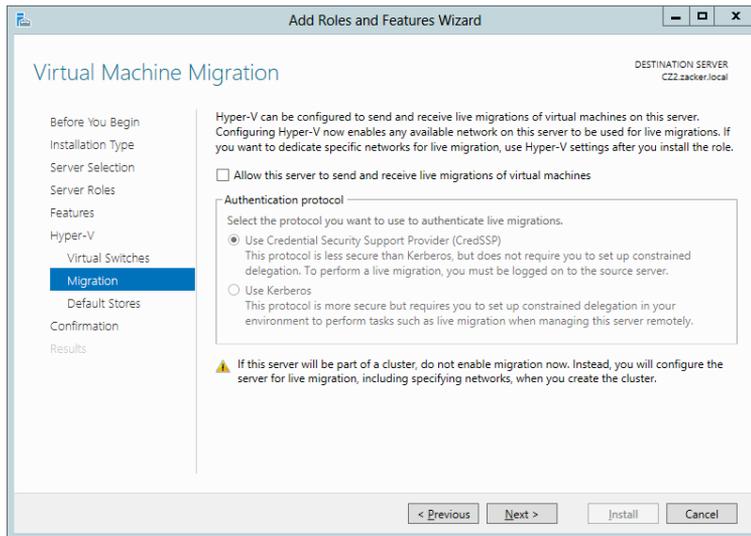


FIGURE 3-6 The Virtual Machine Migration page of the Add Roles and Features Wizard.

11. Click Next to open the Default Stores page.
12. Specify alternatives to the default locations for virtual hard disk (VHD) and VM configuration files, if desired, and click Next. The Confirm Installation Selection page opens.
13. Click Install to move to the Installation Progress page as the wizard installs the role.
14. Click Close to close the wizard.
15. Restart the server.

Installing the role modifies the Windows Server 2012 startup procedure so that the newly installed hypervisor is able to address the system hardware directly and then load the OS as the primary partition on top of that.

NOTE USING WINDOWS POWERSHELL

You can also install the Hyper-V role by using the `Install-WindowsFeature` cmdlet, using the following syntax:

```
Install-WindowsFeature -Name Hyper-V
-ComputerName <name> -IncludeManagementTools -Restart
```

Using Hyper-V Manager

Once you have installed the Hyper-V role and restarted the computer, you can begin to create VMs and deploy OSs on them. The primary graphical tool for creating and managing VMs is the Hyper-V Manager console, which you can access from the Tools menu in Server Manager, just as you can any of the other server and Active Directory management tools.

Like most of the Windows Server 2012 management tools, including Server Manager itself, you can use the Hyper-V Manager console to create and manage VMs on multiple servers, enabling administrators to exercise full control over their servers from a central location.

To run Hyper-V Manager on a server that does not have the Hyper-V role, you must install the Hyper-V Management Tools feature. These tools are also found in the Remote Server Administration Tools package for Windows 8.

Once you install and launch the Hyper-V Manager console, you can add servers to the display by right-clicking the Hyper-V Manager node in the left pane and selecting Connect To Server from the shortcut menu. The Select Computer dialog box appears, in which you can type or browse to the name of a Hyper-V server.

The Hyper-V Manager console lists all the VMs on the selected server, as shown in Figure 3-7, along with status information about each one.

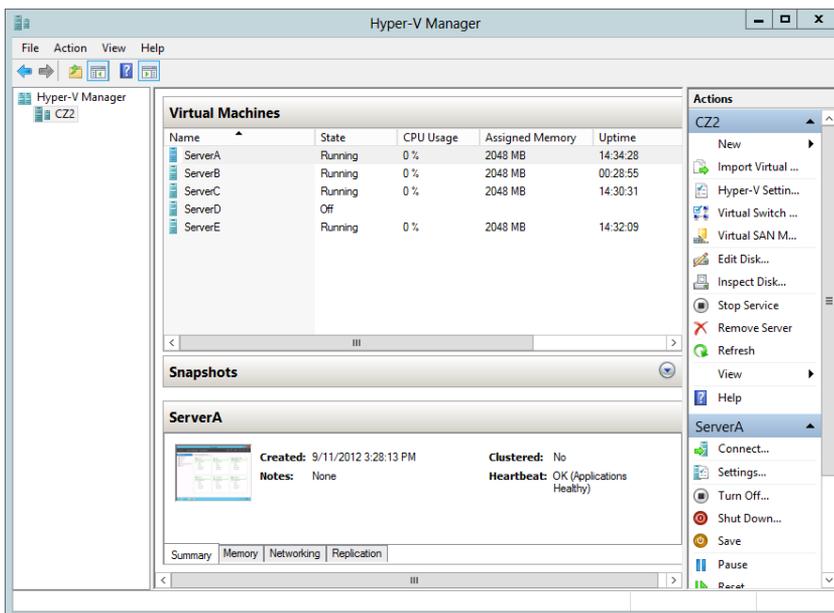


FIGURE 3-7 The Hyper-V Manager console.

Creating a virtual machine

After installing Hyper-V and configuring Hyper-V Manager, you are ready to create VMs and install the OS on each one. By using Hyper-V Manager, you can create new VMs and define the hardware resources that the system should allocate to them. In the settings for a particular VM, depending on the physical hardware available in the computer and the limitations of the guest OS, administrators can specify the number of processors and the amount of memory a VM should use, install virtual network adapters, and create virtual disks by using a variety of technologies, including storage area networks (SANs).

By default, Hyper-V stores the files that make up VMs in the folders you specified on the Default Stores page during the installation. Each VM uses the following files:

- A virtual machine configuration file in XML format with a .xml extension that contains the VM configuration information, including all settings for the VM.
- One or more VHD (.vhd or .vhdx) files to store the guest OS, applications, and data for the VM.

In addition, a VM can use a saved-state (.vsv) file if the machine has been placed into a saved state.

To create a new VM, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager to open the Hyper-V Manager console.
3. In the left pane, select a Hyper-V server.
4. From the Action menu, select New > Virtual Machine. The New Virtual Machine Wizard starts, displaying the Before You Begin page.
5. Click Next to open the Specify Name And Location page.
6. In the Name text box, type a name for the VM, keeping in mind that the system will also use this name to create the VM files and folders. To create the VM files in a location other than the default, select the Store The Virtual Machine In A Different Location check box and type an alternate path in the Location text box. Then click Next. The Assign Memory page opens.

MORE INFORMATION MEMORY

For more information on how Hyper-V uses memory, see the section “Allocating memory” later in this chapter.

7. In the Startup Memory text box, type the amount of memory you want the VM to use and click Next. The Configure Networking page opens, as shown in Figure 3-8.

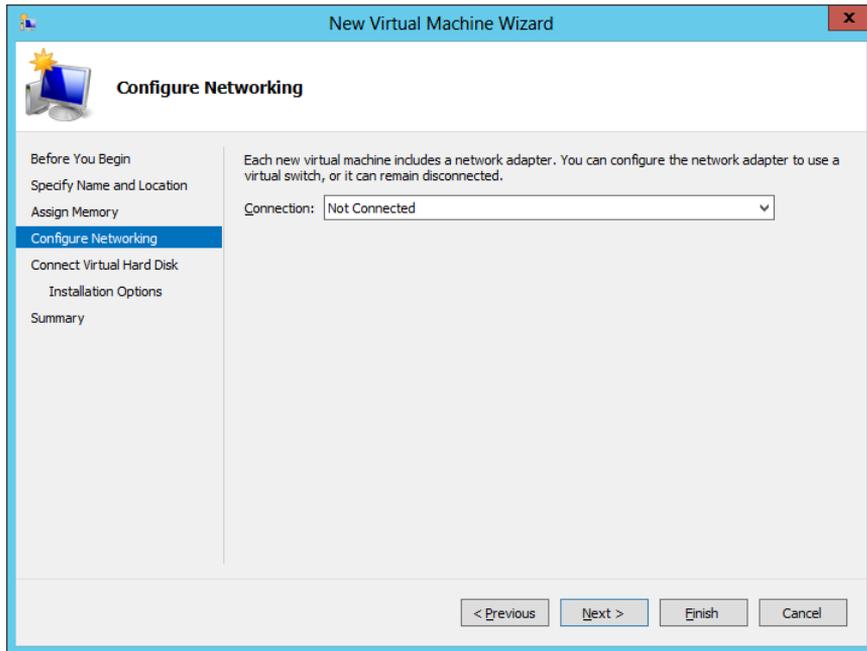


FIGURE 3-8 The Configure Networking page of the New Virtual Machine Wizard.

8. From the Connection drop-down list, select a virtual switch and click Next. The Connect Virtual Hard Disk page opens, as shown in Figure 3-9.

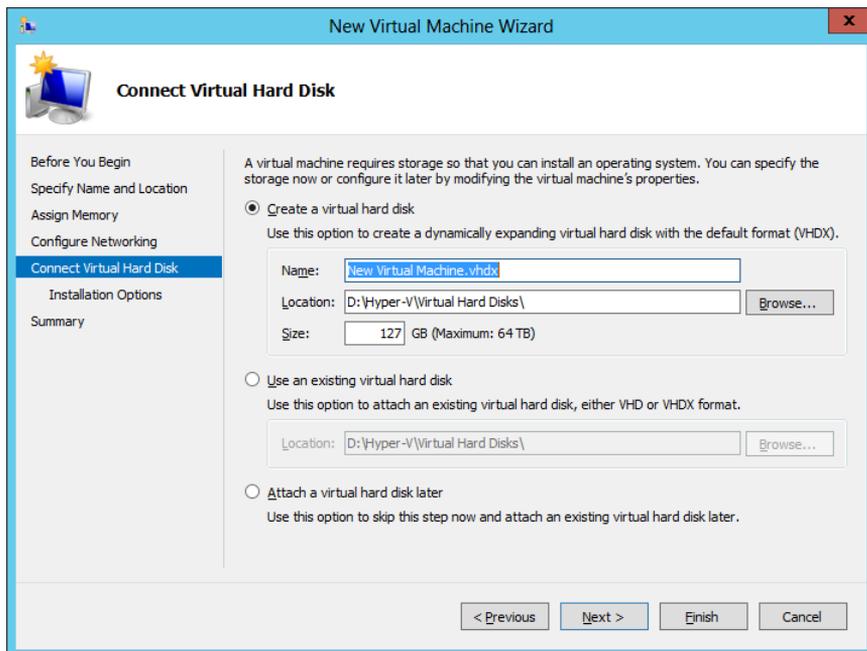


FIGURE 3-9 The Connect Virtual Hard Disk page of the New Virtual Machine Wizard.

MORE INFORMATION NETWORKS

For more information on virtual switches and networking VMs, see Objective 3.3, “Create and configure virtual networks,” later in this chapter.

9. Leave the Create A Virtual Hard Disk option selected and type values for the following fields:
 - **Name** Specifies the file name for the VHD, using the .vhdx format new to Windows Server 2012
 - **Location** Specifies a location for the VHD other than the default you specified on the Default Stores page
 - **Size** Specifies the maximum size of the VHD

MORE INFORMATION STORAGE

By default, the wizard creates a VHD file that starts small and dynamically expands up to the maximum size you specify. For more information on Hyper-V storage, see Objective 3.2, “Create and configure virtual machine storage,” later in this chapter.

10. Click Next. The Installation Options page opens.
11. Leave the Install An Operating System Later Option selected and click Next. The Completing The New Virtual Machine Wizard page opens.
12. Click Finish. The wizard creates the new VM and adds it to the list of VMs in Hyper-V Manager.

The VM that this procedure creates is the equivalent of a bare-metal computer. It has all the (virtual) hardware it needs to run, but it has no software.

NOTE USING WINDOWS POWERSHELL

To create a new VM by using Windows PowerShell, you use the New-VM cmdlet with the following basic syntax:

```
New-VM -Name "VM name" -MemoryStartupBytes <memory>  
  
-NewVHDSIZEBytes <disk size>
```

For example, the following command would create a new VM called ServerA with 1 GB of memory and a new 60-GB VHD drive:

```
New-VM -Name "ServerA" -MemoryStartupBytes 1GB  
  
-NewVHDSIZEBytes 60GB
```

There are, of course, many more parameters for the New-VM cmdlet, which you can explore through the Get-Help cmdlet.

Each VM on a Hyper-V server consists of a collection of settings that specify the hardware resources in the machine and the configuration settings that control those resources. You can manage and modify those settings by using the Settings page for the particular VM.

Selecting a VM from the list in Hyper-V Manager displays a series of icons in the Actions pane. Clicking the Settings icon opens the Settings dialog box, shown in Figure 3-10, which is the primary configuration interface for that VM. Here, you can modify any of the settings that the New Virtual Machine Wizard configured for you.

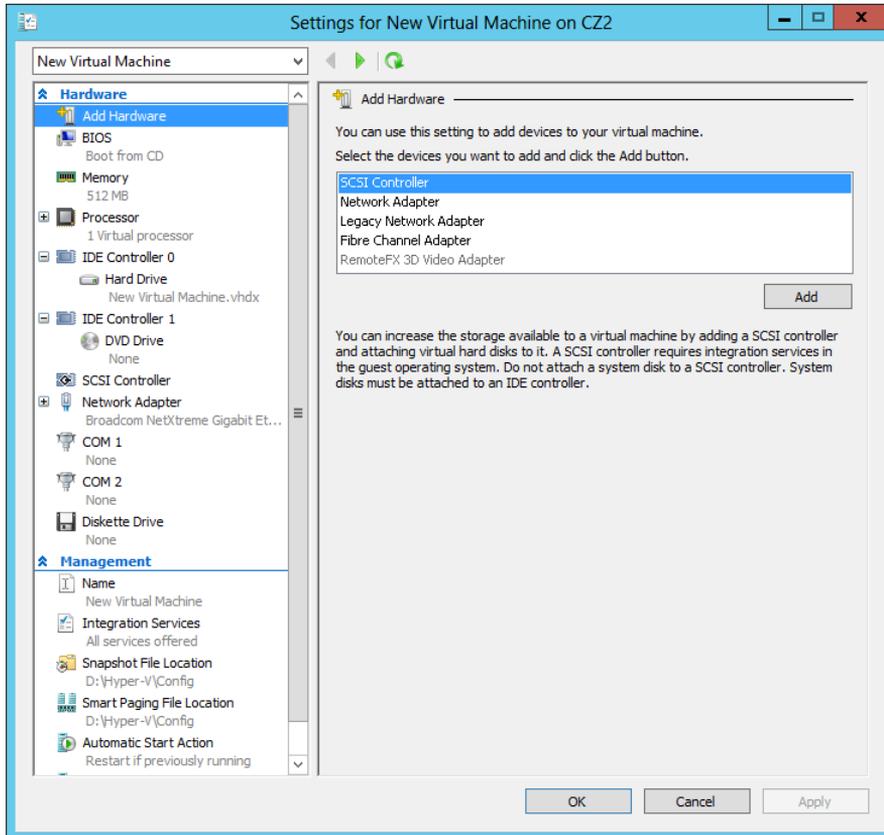


FIGURE 3-10 The Settings dialog box for a VM.

Installing an operating system

Once you have created a VM, you can install an OS on it, just as you would on a new, bare-metal computer. Hyper-V in Windows Server 2012 supports all the following as OSs you can install in VMs:

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

- Windows Home Server 2011
- Windows Small Business Server 2011
- Windows Server 2003 R2
- Windows Server 2003 SP2
- Windows 8
- Windows 7 Enterprise and Ultimate
- Windows Vista Business, Enterprise, and Ultimate SP2
- Windows XP Professional SP3
- Windows XP x64 Professional SP2
- CentOS 6.0–6.2
- Red Hat Enterprise Linux 6.0–6.2
- SUSE Linux Enterprise Server 11 SP2

NOTE GUESTS

This is the official list of supported guest OSs at RTM. Other OSs might also function but have not been fully tested.

One of the advantages of installing software on VMs is that there are several ways to access the installation files. A VM, by default, has a DVD drive, which can itself be physical or virtual.

When you open the Settings dialog box for a VM and select the DVD drive in the Hardware list, you see the interface shown in Figure 3-11. In the Media section, you can select one of the following options for the drive:

- **None** The equivalent of a drive with no disk inserted
- **Image File** Points to a disk image file with a .iso extension stored on one of the host computer's drives or on a shared network drive
- **Physical CD/DVD Drive** Links the virtual DVD drive to one of the physical DVD drives in the host computer

The ability to mount an image file to a virtual DVD drive is particularly useful for administrators who download OS files as disk images. Once you have mounted an installation disk, either physically or virtually, you can click Start in the Actions pane, which is the equivalent of turning on the VM.

Starting a VM causes the thumbnail in the Hyper-V Manager to go live, displaying the contents of the computer's screen. To display the VM's activity at full size, click Connect in the Actions pane to open a new window for the VM. You can then interact with the VM through that window, just as if you were sitting at a physical computer's console.

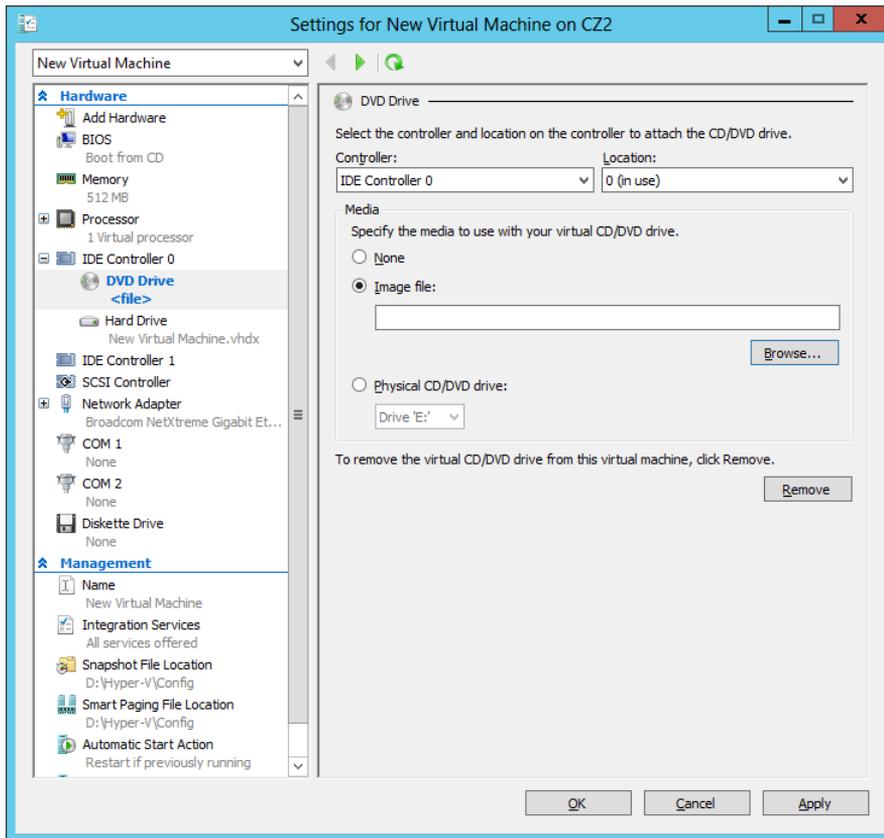


FIGURE 3-11 DVD drive settings for a VM.

When the VM boots from the disk you mounted, the OS installation proceeds just as if you were using a physical computer. During the installation process, you can work with the VHD drive just as you would a physical one, creating partitions of various sizes and selecting one for the OS. When the installation is complete, the VM restarts, and you can then log on and use it in the normal manner.

Configuring Guest Integration Services

In some cases, certain Hyper-V guest OS features do not function properly using the OS's own device drivers. Hyper-V, therefore, includes a software package called Guest Integration Services, which you can install on your VMs for compatibility purposes.

Some of the functions provided by the Guest Integration Services package are as follows:

- **Operating system shutdown** Enables the Hyper-V Manager console to remotely shut down a guest OS in a controlled manner, eliminating the need for an administrator to log on and manually shut the system down.

- **Time synchronization** Enables Hyper-V to synchronize the OS clocks in parent and child partitions.
- **Data Exchange** Enables the Windows OSs on the parent and child partitions to exchange information, such as OS version information and fully qualified domain names.
- **Heartbeat** Implements a service in which the parent partition sends regular heart-beat signals to the child partitions, which are expected to respond in kind. A failure of a child partition to respond indicates that the guest OS has frozen or malfunctioned.
- **Backup** Enables backup of Windows VMs by using Volume Shadow Copy Services.

The Windows Server 2012 and Windows 8 operating systems have the latest Guest Integration Services software built in, so there is no need to install the package on VMs running those OSs as guests. Earlier versions of Windows have earlier versions of the Guest Integration Services package that need to be upgraded, however, and some Windows versions do not include the package at all.

NOTE LINUX

For Linux guest OSs, you must download and install the latest release of Linux Integration Services Version 3.4 for Hyper-V from the Microsoft Download Center. As of this writing, the latest version is 3.4 and is available at <http://www.microsoft.com/en-gb/download/details.aspx?id=28188>.

To upgrade Guest Integration Services on a Windows guest OS, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager. The Hyper-V Manager console starts.
3. In the left pane, select a Hyper-V server.
4. In the Actions pane, start the VM on which you want to install Guest Integration Services and click Connect. A Virtual Machine Connection window opens.
5. In the Virtual Machine Connection window, from the Action menu, select Insert Integration Services Setup Disk. Hyper-V mounts an image of the Guest Integration Services disk to a virtual disk drive, and an Autoplay window appears.
6. Click Install Hyper-V Integration Services. A message box appears, asking you to upgrade the existing installation.
7. Click OK. The system installs the package and prompts you to restart the computer.
8. Click Yes to restart the computer.

Once you have installed or upgraded Guest Integration Services, you can enable or disable each of the individual functions by opening the Settings dialog box for the VM and selecting the Integration Services page, as shown in Figure 3-12.

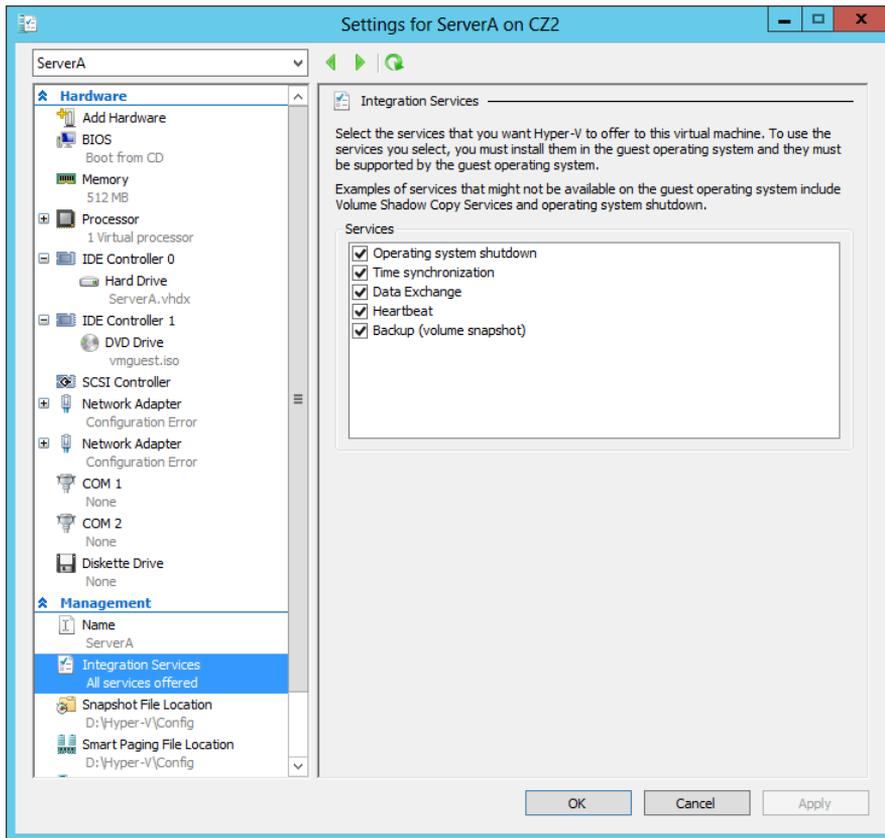


FIGURE 3-12 Integration Services settings for a VM.

At this point, you are ready to configure and manage the VM just as if you were working on a physical server. This can include modifying the network configuration, enabling remote desktop, loading the appropriate roles and features, and installing applications.

Allocating memory

Dynamic memory enables Hyper-V to adjust the amount of RAM allocated to VMs, depending on their ongoing requirements. Some computer components can be virtualized. You can take some disk space and create a virtual hard drive, and you can take an image file and create a virtual DVD drive. You can also create virtual network interface adapters and other components, which appear like the real thing in a VM. System memory is different, however. There is no substitute for memory, so all Hyper-V can do is take the physical memory installed in the computer and allocate it among the various VMs.

When you create a VM by using the New Virtual Machine Wizard, you specify how much memory the VM should use on the Assign Memory page. Obviously, the amount of memory available for use is based on the physical memory installed in the computer.

After you have created the VM, you can modify the amount of memory allocated to it by shutting down the VM, opening its Settings dialog box, and changing the Startup RAM setting on the Memory page, as shown in Figure 3-13. This enables you to experiment with various amounts of memory, and set the optimum performance level for the system.

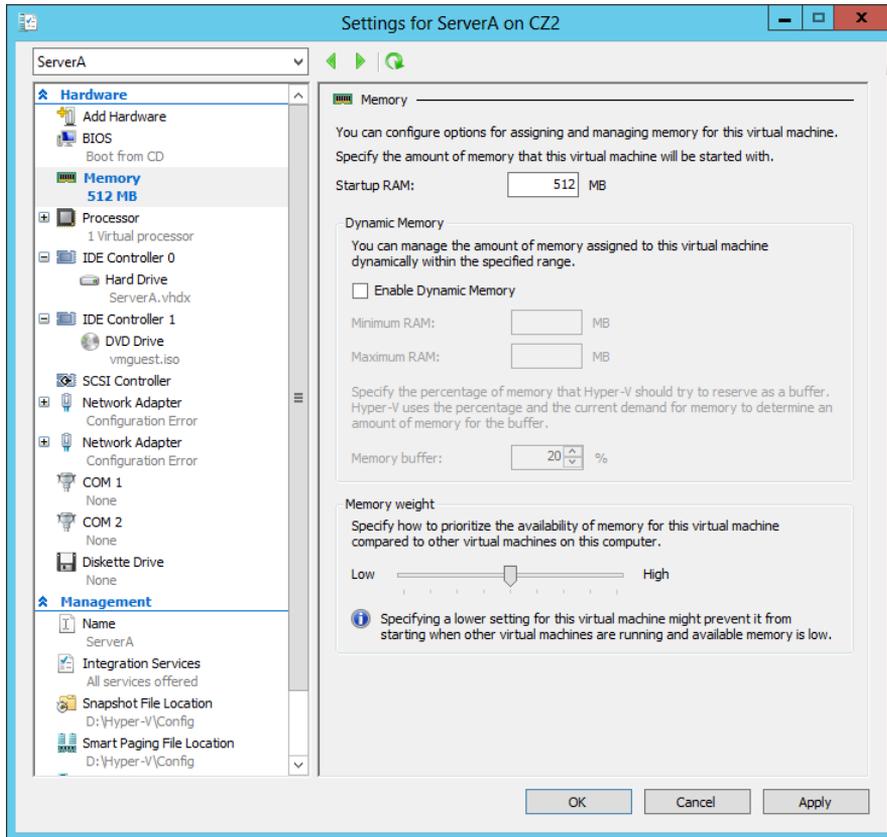


FIGURE 3-13 Memory settings for a VM.

USING DYNAMIC MEMORY

In the first versions of Hyper-V, shutting down the VM was the only way to modify its memory allocation. In the Windows Server 2012 version, however, you can use a feature called Dynamic Memory to automatically reallocate memory to the VM from a shared memory pool as its demands change. If a virtualized server starts to experience larger amounts of client traffic, for example, Hyper-V can increase the memory allocated to the system, and reduce it again when the traffic subsides.

To use Dynamic Memory, you must enable it by selecting the Enable Dynamic Memory check box on the VM's Memory settings page, and then configure the following settings:

- **Startup RAM** Specifies the amount of memory that you want to allocate to the VM when it starts. When you are using Dynamic Memory, this value can be the minimum amount of memory needed to boot the system.
- **Minimum RAM** Specifies the smallest amount of memory the VM can use at any time. OSs can require more memory to start up than to run, so this value can be smaller than the Startup RAM value.
- **Maximum RAM** Specifies the largest amount of memory that the VM can use at any time. The value can range from a low equal to the Startup RAM value to a high of 64 GB.
- **Memory Buffer** Specifies a percentage that Hyper-V uses to calculate how much memory to allocate to the VM, compared to its actual utilization, as measured by performance counters. For example, with the Memory Buffer value set to 20 percent, a VM with applications and OS that consume 1 GB of memory will receive a dynamic allocation of 1.2 GB.
- **Memory Weight** Specifies a relative value that specifies the priority of this VM, compared to the other VMs on the same computer. When the physical memory in the computer is insufficient to allocate the full buffered amount specified for each VM, the VMs with the highest Memory Weight settings receive priority.

NOTE RAM

You can reduce the Minimum RAM, increase the Maximum RAM, or change the Memory Buffer and Memory Weight values at any time, but to enable or disable Dynamic Memory, you must shut down the VM.

In addition to configuring the VM settings, the guest VM must be running Windows Vista or later or Windows Server 2003 SP2 or later and have Windows Server 2012 Guest Integration Services installed to use Dynamic Memory.

NOTE USING WINDOWS POWERSHELL

To configure the memory settings for a VM, you use the Set-VMMemory cmdlet, using the following basic syntax:

```
Set-VMMemory <VM name> -DynamicMemoryEnabled $true
-MinimumBytes <memory> -StartupBytes <memory>
-MaximumBytes <memory> -Priority <value> -Buffer <percentage>
```

For example, to configure the memory settings for the VM ServerA, enabling Dynamic Memory and configuring values for all of its settings, use the following command:

```
Set-VMMemory TestVM -DynamicMemoryEnabled $true
-MinimumBytes 64MB
```

CONFIGURING SMART PAGING

Dynamic Memory was introduced in Windows Server 2008 R2 Hyper-V, but Windows Server 2012 improves on the concept by adding the Minimum RAM setting. This makes it possible for Hyper-V to reduce the memory used by a VM to a level lower than that needed to start the system, reclaiming that memory for other uses.

The problem with having minimum RAM values that are lower than the startup RAM values is that it becomes possible to deplete the supply of physical memory with too many VMs running simultaneously at their minimum RAM values. If this occurs, a VM that has to restart might be unable to do so because there is not enough free memory to increase its memory allocation from its minimum RAM value to its startup RAM value.

To address this possibility, Hyper-V includes a feature called smart paging. If a VM has to restart and there is not enough memory available to allocate its startup RAM value, the system uses hard disk space to make up the difference and begins paging memory contents to disk.

Disk access rates are far slower than memory access rates, of course, so smart paging incurs a severe performance penalty, but the paging occurs only for as long as it takes to restart the VM and return it to its minimum RAM allocation.

Hyper-V only uses smart paging in specific conditions: when a VM must be restarted, there is no free memory available, and there are no other means available to free up the necessary memory.

You can select the Smart Paging File Location page in a VM's Setting dialog box to specify a location for the paging file. Selecting the fastest possible hard drive is recommended.

Configuring resource metering

Resource metering is a new Windows PowerShell–based feature in Windows Server 2012 Hyper-V that enables administrators to document VM usage by using a variety of criteria. There are various reasons why organizations might want to track the use of VMs. For large corporations, it might be a matter of internal accounting and controlling ongoing expenses, such as wide area network (WAN) bandwidth. For service providers, it might be necessary to bill customers based on the VM resources they use.

Resource metering uses Windows PowerShell cmdlets to track a variety of performance metrics for individual VMs, including the following:

- CPU utilization
- Minimum, maximum, and average memory utilization
- Disk space utilization
- Incoming and outgoing network traffic

Resource metering statistics remain consistent, even when you transfer VMs between host systems by using Live Migration or move VHD files between VMs.

To use resource metering, you must first enable it for the specific VM that you want to monitor by using the `Enable-VMResourceMetering` cmdlet with the following syntax:

```
Enable-VMResourceMetering -VMName <name>
```

Once you have enabled metering, you can display a statistical report at any time by using the `Measure-VM` cmdlet with the following syntax:

```
Measure-VM -VMName <name>
```

In addition to metering resources for entire VMs, administrators can also create resource pools that enable them to monitor specific VM components, such as processors, memory, network adapters, and VHDs. You create a resource pool by using the `New-VMResourcePool` cmdlet and then enable metering for the pool by using `Enable-VMResourceMetering`.

By using techniques such as pipelining, administrators can use the resource metering cmdlets to gather data on VM performance and export it to applications or data files.

Objective summary

- Virtualization is a process that adds a layer of abstraction between actual, physical hardware and the system making use of it. Instead of having the server access the computer's hardware directly, an intervening component called a hypervisor creates a VM environment, and the server OS runs in that environment.
- Virtualization is the process of deploying and maintaining multiple instances of an OS, called VMs, on a single computer.
- Microsoft Hyper-V is a hypervisor-based virtualization system for x64 computers starting with Windows Server 2008. The hypervisor is installed between the hardware and the OS and is the main component that manages the virtual computers.
- For licensing purposes, Microsoft refers to each VM that you create on a Hyper-V server as a virtual instance. Each Windows Server 2012 version includes a set number of virtual instances; you must purchase licenses to create additional instances.
- To keep a small footprint and minimal overhead, Hyper-V Server contains only the Windows Hypervisor, Windows Server driver model, and virtualization components.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of this chapter.

1. Which of the following statements about Type I and Type II virtualization are true? (Choose all that apply.)
 - A. In Type I virtualization, the hypervisor runs on top of a host OS.
 - B. In Type I virtualization, the hypervisor runs directly on the computer hardware.

- C.** In Type II virtualization, the hypervisor runs on top of a host OS.
 - D.** In Type II virtualization, the hypervisor runs directly on the computer hardware.
- 2.** Which of the following types of server virtualization provides the best performance for high-traffic servers in production environments?
 - A.** Type I virtualization
 - B.** Type II virtualization
 - C.** Presentation virtualization
 - D.** RemoteApp
- 3.** Which of the following Microsoft operating systems includes a license that enables you to create an unlimited number of virtual instances?
 - A.** Hyper-V Server
 - B.** Windows Server 2012 Datacenter
 - C.** Windows Server 2012 Standard
 - D.** Windows Server 2012 Foundation
- 4.** Which of the following Hyper-V features make it possible for a VM to function with a minimum RAM value that is lower than the startup RAM value? (Choose all that apply.)
 - A.** Smart paging
 - B.** Dynamic Memory
 - C.** Memory Weight
 - D.** Guest Integration Services
- 5.** When you install the Hyper-V role on a server running Windows Server 2012, the instance of the OS on which you installed the role is converted to what system element?
 - A.** The hypervisor
 - B.** The Virtual Machine Monitor
 - C.** The parent partition
 - D.** A child partition



Thought experiment

In the following thought experiment, apply what you've learned about this objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

Alice has a computer running Windows Server 2012 with 8 GB of memory installed, which she has configured as a Hyper-V server. After creating eight VMs by using the New Virtual Machine Wizard, each with a startup RAM value of 1,024 MB, Alice is having trouble getting all eight VMs to boot. What settings can she modify to resolve the problem without changing the startup RAM value?

Objective 3.2: Create and configure virtual machine storage

When you create a VM in Windows Server 2012 Hyper-V, you emulate all the standard components that you typically find in a physical computer. When you virtualize memory, as discussed in Objective 3.1, "Create and configure virtual machine settings," you take a portion of the physical memory in the computer and dedicate it to a VM. The same is true with hard disk space. Hyper-V uses a specialized VHD format to package part of the space on a physical disk and make it appear to the VM as though it is a physical hard disk drive.

When you create a new VM in Hyper-V by using the New Virtual Machine Wizard, the wizard creates a virtual storage subsystem that consists of two Integrated Drive Electronics (IDE) controllers and one Small Computer Systems Interface (SCSI) controller. The IDE controllers host the VM's system drive and its DVD drive. Like their physical equivalents, each IDE controller can host two devices, so you can create two additional virtual drives and add them to the system.

The SCSI controller in the default VM configuration is unpopulated, and you can create additional drives and add them to that controller to provide the VM with additional storage. You can also create additional SCSI controllers and add drives to them. By creating multiple drives and controllers, Hyper-V makes it possible to construct virtual storage subsystems that emulate almost any physical storage solution you might devise.

This objective covers how to:

- Create VHDs and VHDX
- Configure differencing drives
- Modify VHDs
- Configure pass-through disks
- Manage snapshots
- Implement a virtual Fibre Channel adapter

Virtual disk formats

Windows Server 2012 Hyper-V supports the original VHD disk image file and the new VHDX format. The original VHD format was created by a company called Connectix for its Virtual PC product. Microsoft later acquired the product and used the VHD format for all its subsequent virtualization products, including Hyper-V. There are three types of VHD files, as follows:

- **Fixed hard disk image** An image file of a specified size in which all the disk space required to create the image is allocated during its creation. Fixed disk images can be wasteful in terms of storage because they can contain large amounts of empty space, but they are also efficient from a processing standpoint because there is no overhead due to dynamic expansion.
- **Dynamic hard disk image** An image file with a specified maximum size, which starts small and expands as needed to accommodate the data the system writes to it.
- **Differencing hard disk image** A child image file associated with a specific parent image. The system writes all changes made to the data on the parent image file to the child image, to manage disk space or to facilitate a rollback at a later time.

VHD images are limited to maximum size of 2 TB and are compatible with all versions of Hyper-V and Microsoft Type II hypervisor products, such as Virtual Server and Virtual PC. Windows Server 2012 introduced an updated version of the format, which uses a VHDX file-name extension.

VHDX image files can be as large as 64 TB, and they also support 4-KB logical sector sizes to provide compatibility with new 4-KB native drives. VHDX files can also use larger block sizes (up to 256 MB), which enable administrators to fine-tune the performance level of a virtual storage subsystem to accommodate specific applications and data file types. However, VHDX files are not backward compatible and can only be read by Windows Server 2012 and Windows 8 Hyper-V servers. If migrating your VMs from Windows Server 2012 to an older version of Hyper-V is even a remote possibility, you should continue using the VHD file format.

Creating virtual disks

Windows Server 2012 Hyper-V provides several ways to create virtual disk files. You can create them as part of a VM or create them at another time and add them to a VM. The graphical interface in Hyper-V Manager provides access to most of the VHD parameters, but the new Windows PowerShell cmdlets included in Windows Server 2012 provide the most granular control over the disk image format.

Creating a virtual disk with a VM

The New Virtual Machine Wizard includes a Connect Virtual Hard Disk page with which you can add a single disk to your new VM. The options for this disk are relatively limited and consist of the following:

- **Create A Virtual Hard Disk** Enables you to specify the name, location, and size of a new VHD. You can only create a dynamically expanding disk using the VHDX format.
- **Use An Existing Virtual Hard Disk** Enables you to specify the location of an existing VHD or VHDX disk, which the VM will presumably use as its system disk.
- **Attach A Virtual Hard Disk Later** Prevents the wizard from adding any virtual disks to the VM configuration. The assumption is that you will manually add a disk later, before you start the VM.

The object of this wizard page is to create the disk on which you will install the VM's OS or to select an existing disk on which an OS is already installed. The disk the wizard creates is always a dynamically expanding one connected to IDE Controller 0.

NOTE VHDS

It has become a common practice for Microsoft to release evaluation copies of its products as preinstalled VHD files, as an alternative to the traditional installable disk images. After downloading one of these files, you can create a VM on a Hyper-V server and select the Use An Existing Virtual Hard Disk option to mount the VHD as its system drive.

Creating a new virtual disk

You can create a VHD file at any time, without adding it to a VM, by using the New Virtual Hard Disk Wizard in Hyper-V Manager. To create a new virtual disk, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager. The Hyper-V Manager console opens.
3. In the left pane, select a Hyper-V server.
4. From the Action menu, select New > Hard Disk to start the New Virtual Hard Disk Wizard, displaying the Before You Begin page.

5. Click Next to open the Choose Disk Format page.
6. Select one of the following disk format options:
 - **VHD** Creates an image no larger than 2 TB, using the highly compatible VHD format
 - **VHDX** Creates an image up to 64 TB, using the new VHDX format
7. Click Next to open the Choose Disk Type page.
8. Select one of the following disk type options:
 - **Fixed Size** Creates a disk of a specific size, allocating all of the space at once
 - **Dynamically Expanding** Creates a disk that grows to the maximum size you specify as you add data
 - **Differencing** Creates a child drive that will contain changes made to a specified parent drive
9. Click Next. The Specify Name And Location page opens.
10. Specify a file name for the disk image in the Name text box and, if desired, specify a location for the file other than the server default. Click Next to open the Configure Disk page.
11. For fixed and dynamically expanding disks, select and configure one of the following options:
 - **Create A New Blank Virtual Hard Disk** Specifies the size (or the maximum size) of the disk image file to create
 - **Copy The Contents Of The Specified Physical Disk** Enables you to select one of the physical hard disks in the computer and copy its contents to the new disk image
 - **Copy The Contents Of The Specified Virtual Hard Disk** Enables you to select an existing virtual disk file and copy its contents to the new disk image
12. Click Next. The Completing The New Virtual Hard Disk Wizard page opens.
13. Click Finish.

The wizard creates the new image disk and saves it to the specified location.

NOTE USING WINDOWS POWERSHELL

You can create new VHD files by using Windows PowerShell, which gives you more control than is available through the graphical interface. To create a new disk image, use the New-VHD cmdlet with the following basic syntax:

```
New-VHD -Path c:\filename.vhd|c:\filename.vhdx
-Fixed|-Dynamic|-Differencing -SizeBytes <size>
[-BlockSizeBytes <block size>]
[-LogicalSectorSizeBytes 512|4096] [-ParentPath <pathname>]
```

When using the cmdlet to create a disk image, the extension you specify for the filename determines the format (VHD or VHDX), and you can specify the block size and the logical

sector size for the image, two things you cannot do in the GUI. For example, the following command will create a 400-GB fixed VHDX image file with a logical sector size of 4 KB:

```
New-VHD -Path c:\diskfile.vhdx -Fixed  
-SizeBytes 400GB -LogicalSectorSizeBytes 4096
```

Adding virtual disks to virtual machines

Creating virtual disk image files as a separate process enables administrators to exercise more control over their capabilities, but after creating the VHD or VHDX files, you must add them to a VM for them to be useful.

To add a hard disk drive to a physical computer, you must connect it to a controller, and the same is true with a VM in Hyper-V. When you open the Settings dialog box for a VM in its default configuration, you see three controllers, labeled IDE Controller 0, IDE Controller 1, and SCSI Controller. These correspond to the controllers you might find in a typical physical server computer.

Each of the IDE controllers can support two devices, and the default VM configuration uses one channel on IDE Controller 0 for the system hard disk and one channel on IDE controller 1 for the system's DVD drive. If you did not create a virtual disk as part of the new Virtual Machine Wizard—that is, if you chose the Attach A Virtual Hard Disk Later option—then you must add a hard disk image to IDE Controller 0 to use as a system drive. The VM cannot boot from the SCSI controller.

To add an existing virtual system drive to a VM, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager to open the Hyper-V Manager console.
3. In the left pane, select a Hyper-V server.
4. Select a VM and, in the Actions pane, select Settings. The Settings dialog box for the VM appears.
5. Select IDE Controller 0, as shown in Figure 3-14.

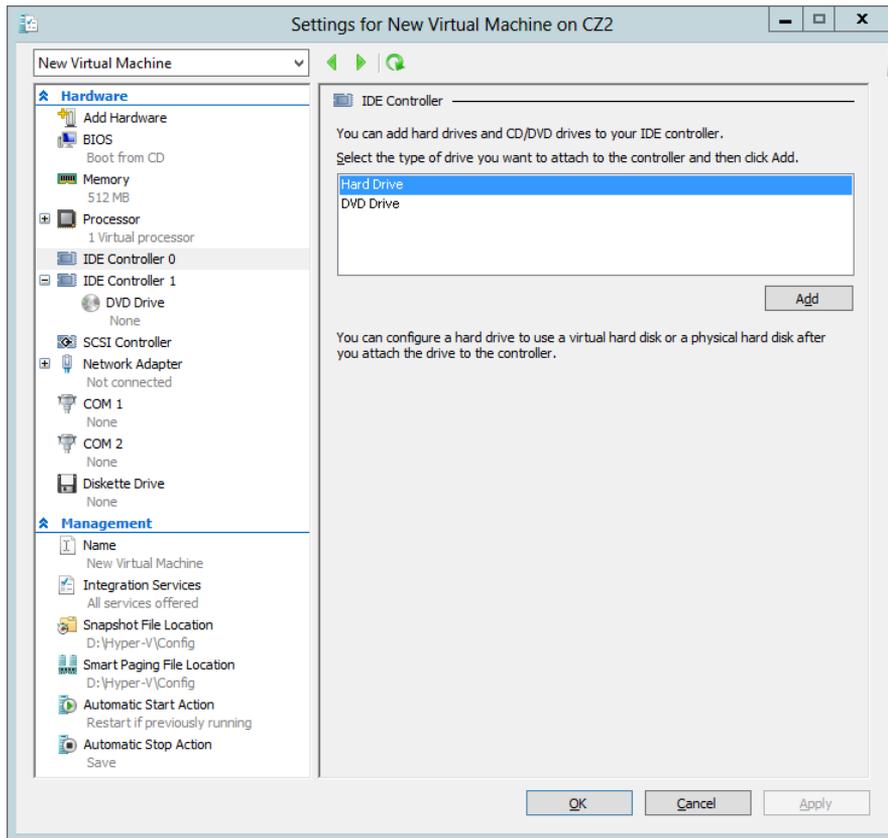


FIGURE 3-14 The IDE Controller interface in the Settings dialog box.

6. In the IDE Controller box, select Hard Drive and click Add. The Hard Drive page opens, as shown in Figure 3-15.

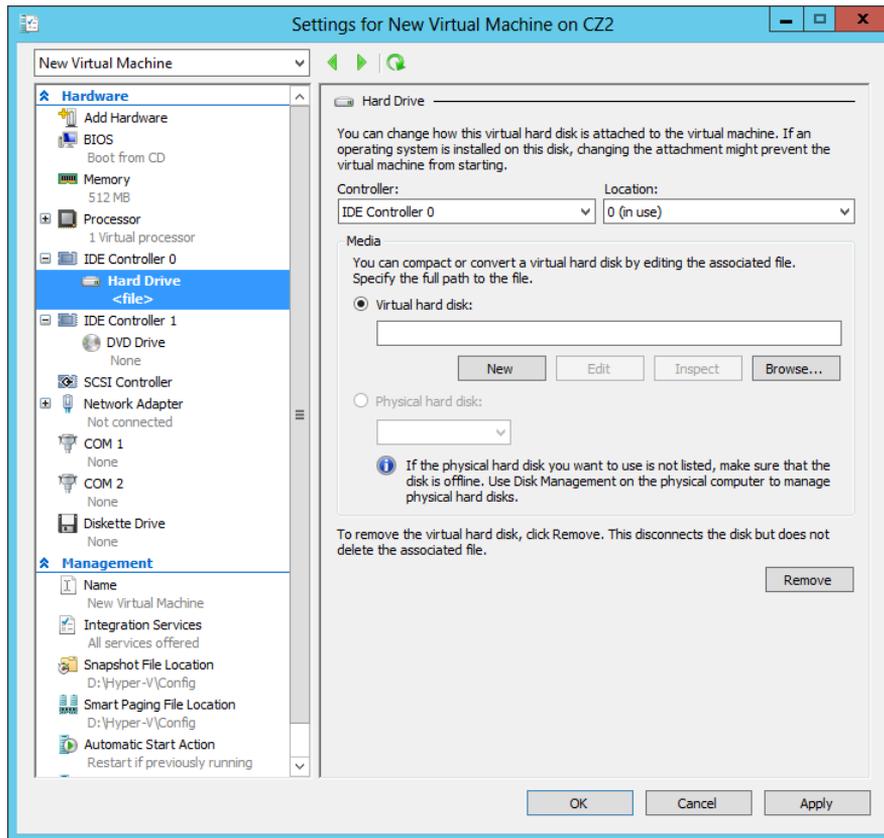


FIGURE 3-15 The Hard Drive interface in the Settings dialog box.

7. In the Controller and Location drop-down lists, select the IDE controller and the channel you want to use for the hard disk.
8. With the Virtual Hard Disk option selected, click Browse and select the disk image file you want to add.
9. Click OK to close the Settings dialog box.

Although you cannot use a SCSI drive as the system disk in a VM, you can add virtual data disks to the SCSI controller. Unlike the IDE connectors, which support only two devices each, a SCSI connector in Hyper-V can support up to 64 drives. You can also add SCSI controllers to a VM, providing almost unlimited scalability for your virtual storage subsystem.

Creating differencing disks

A differencing disk enables you to preserve an existing virtual disk image file in its original state while mounting it in an operating system and even modifying its contents. For example, when building a laboratory setup, you can create a baseline system by installing a clean copy of an OS on a new virtual disk and configuring the environment to fit your needs. Then, you

can create a new child differencing disk, using your baseline image as the parent. All subsequent changes you make to the system will then be written to the differencing disk while the parent remains untouched. You can experiment on the test system as you wish, knowing that you can revert to your baseline configuration by just creating a new differencing disk.

You can create multiple differencing disks that point to the same parent image, enabling you to populate a lab network with as many VMs as you need, which saves disk space and eliminates the need to repeatedly install the OS.

To create a cloned version of a baseline installation with a differencing disk, use the following procedure.

- 1.** Install and configure the baseline VM. Create a new VM with a new disk image file and install a guest OS on it. Configure the OS as needed and install any roles, features, applications, or services you need.
- 2.** Generalize the parent image. Open an elevated command prompt on the baseline system and run the Sysprep.exe utility. Sysprep configures the system to assign itself a new, unique security ID (SID) the next time the computer starts. This enables you to create multiple cloned systems from a single disk image.
- 3.** Create a parent disk image. Once you have generalized the baseline installation, you no longer need the original VM. You can delete everything except the VHD or VHDX file containing the disk image. This will become your parent image. Open the Properties sheet for the image file and set the read-only flag, to ensure that the baseline does not change.
- 4.** Create a differencing disk. By using the New Virtual Hard Disk Wizard or the New-VHD cmdlet for Windows PowerShell, create a new differencing disk pointing to the baseline image you created and prepared earlier as the parent image.
- 5.** Create a cloned VM. Create a new VM and, on the Connect Virtual Hard Disk page, attach the differencing disk you just created to it by using the Use An Existing Virtual Hard Disk option.

You can then proceed to create additional cloned VMs with differencing disks that all use the same parent. Each one can function independently, and the parent disk will remain unchanged.

When you create a differencing drive by using the New Virtual Hard Disk Wizard, selecting the Differencing option on the Choose Disk Type page causes the Configure Disk page to appear as shown in Figure 3-16. In the Location text box, specify the name of the file that you want to use as the parent image.

In the same way, if you create the differencing disk by using Windows PowerShell, you must run the New-VHD cmdlet with the `-Differencing` parameter and the `-ParentPath` parameter, specifying the location of the parent disk.

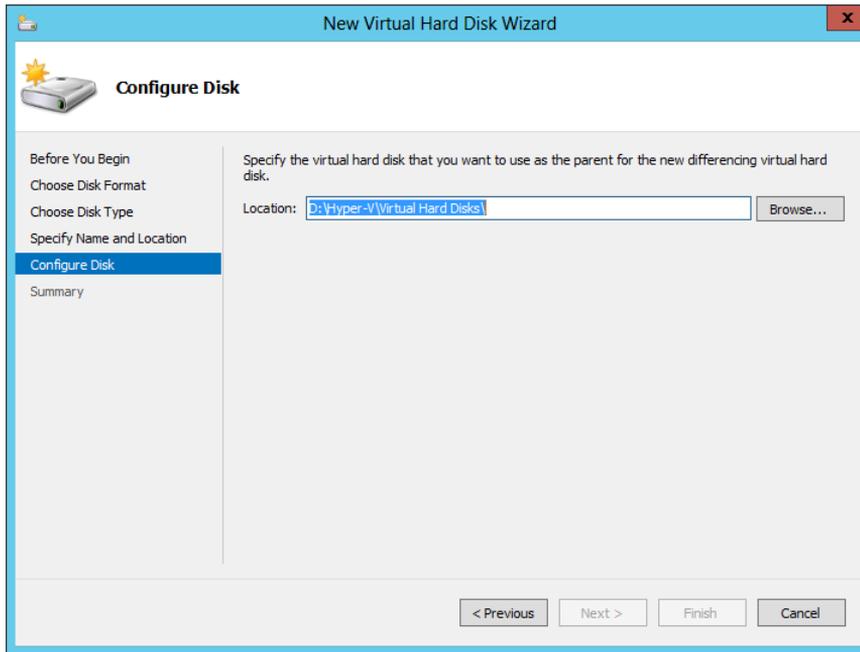


FIGURE 3-16 The Configure Disk page in the New Virtual Hard Disk Wizard when creating a differencing disk.

Configuring pass-through disks

This objective has thus far been concerned primarily with VHDs, areas of space on a physical disk drive allocated for use by VMs. However, it is also possible for VMs to access physical disks directly.

A pass-through disk is a type of virtual disk that points not to an area of space on a physical disk, but to a physical disk drive installed on the host computer. When you add a hard drive to any of the controllers in a VM, you have the option of selecting a physical hard disk as opposed to a virtual one.

To add a physical hard disk to a VM, the VM must have exclusive access to it. This means that you must first take the disk offline in the parent OS by using the Disk Management snap-in, as shown in Figure 3-17, or the Diskpart.exe utility. Once the disk is offline, it will be available for selection in the Physical Hard Disk drop-down list.

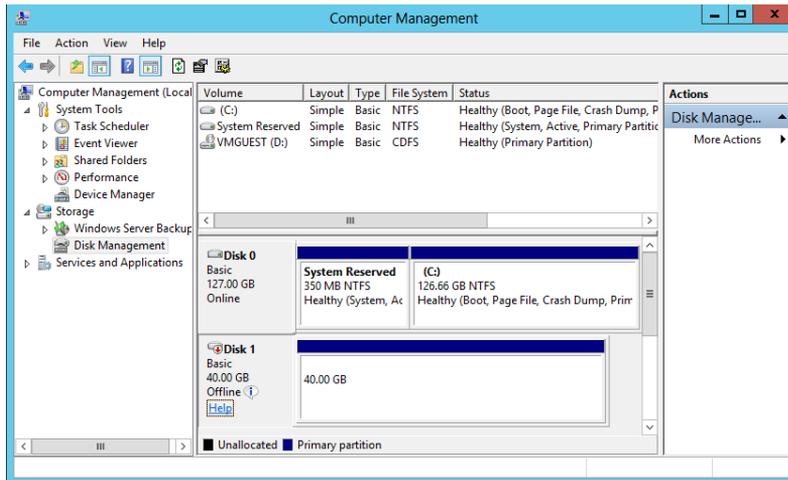


FIGURE 3-17 An offline disk in the Disk Management snap-in.

Modifying virtual disks

Windows Server 2012 and Hyper-V provide several ways for administrators to manage and manipulate VHD images without mounting them in a VM. Once you have created a VHD, whether you have attached it to a VM or not, you can manage it by using the Edit Virtual Hard Disk Wizard in Hyper-V Manager. To edit an existing VHD or VHDX file, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager to open the Hyper-V Manager console.
3. In the left pane, select a Hyper-V server.
4. In the Actions pane, select Edit Disk. The Edit Virtual Hard Disk Wizard starts, displaying the Before You Begin page.
5. Click Next to open the Locate Disk page.
6. Type or browse to the name of the VHD or VHDX file you want to open and click Next. The Choose Action page appears.
7. Select one of the following functions:
 - **Compact** Reduces the size of a dynamically expanding or differencing disk by deleting empty space while leaving the disk's capacity unchanged
 - **Convert** Changes the type of format of a disk by copying the data to a new disk image file
 - **Expand** Increases the capacity of the disk by adding empty storage space to the image file

- **Shrink** Reduces the capacity of the disk by deleting empty storage space from the file
 - **Merge** Combines the data on a differencing disk with that of the parent disk to form a single composite image file
8. Click Next to open the Completing The Edit Virtual Hard Disk Wizard page.
 9. Complete any new pages presented by the wizard as a result of your selection and click Finish.

The options that appear on the wizard's Choose Action page depend on the current status of the image file you select. For example, the Merge option only appears if you choose a differencing disk, and the Shrink option does not appear unless there is free space in the file that the wizard can delete.

In addition to these disk editing functions provided by Hyper-V Manager, it is possible to use the Disk Management snap-in on the Hyper-V host to mount a VHD or VHDX file as a drive and access its contents, just as if it were a physical disk.

To mount a VHD file, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Computer Management to open the Computer Management console.
3. In the left pane, select Disk Management. The Disk Management snap-in opens.
4. From the Action menu, select Attach VHD. The Attach Virtual Hard Disk dialog box appears.
5. In the Location text box, type or browse to the image disk file you want to attach and click OK. The disk appears in the Disk Management interface.
6. Close the Computer Management console.

At this point, you can work with the virtual disk and its contents using any standard tools, just as you would a physical hard disk drive. To detach the VHD, you use the same procedure and select Detach VHD from the Action menu.

Creating snapshots

In Hyper-V, a snapshot is a captured image of the state, data, and hardware configuration of a VM at a particular moment in time. Creating snapshots is a convenient way for administrators to revert a VM to a previous state at will. For example, if you create a snapshot just before applying a system update, and the update is somehow problematic, you can apply the snapshot and return the VM to the state in which it was before you applied the update.

Creating a snapshot is as simple as selecting a running VM in Hyper-V Manager and selecting Snapshot from the Actions pane. The system creates a snapshot file, with an AVHD or AVHDX extension, in the same folder as the VHD file, and adds the snapshot to the Hyper-V Manager display, as shown in Figure 3-18.

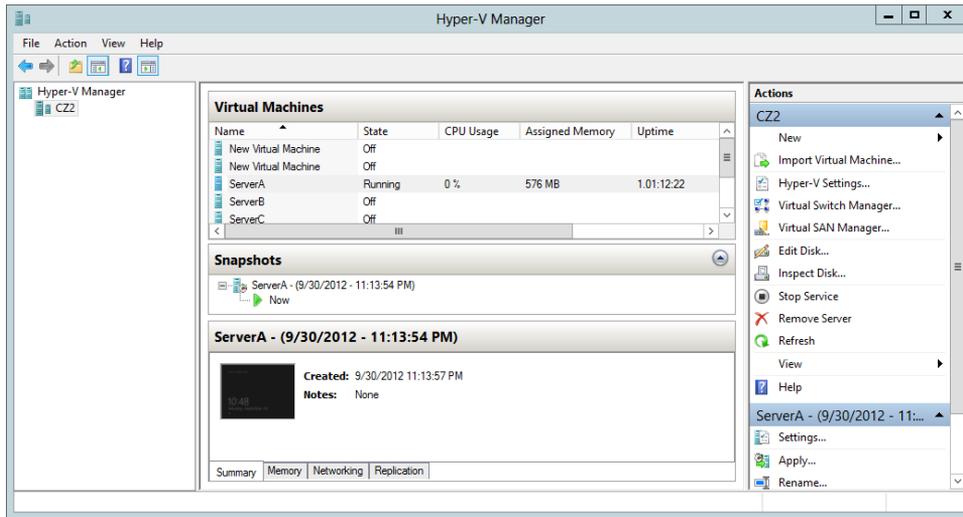


FIGURE 3-18 A snapshot in Hyper-V Manager.

Snapshots are a useful tool for administrators implementing a test environment in Hyper-V, but they are not recommended for heavy use in production environments. In addition to consuming disk space, the presence of snapshots can reduce the overall performance of a VM's disk subsystem.

Connecting to a SAN

At its most basic level, a SAN is simply a network dedicated to high-speed connections between servers and storage devices. Instead of installing disk drives into servers or connecting them by using an external SCSI bus, a SAN consists of one or more drive arrays equipped with network interface adapters, which you connect to your servers by using standard twisted pair or fiber optic network cables. A SAN-connected server, therefore, has a minimum of two network adapters, one for the standard local area network (LAN) connection and one for the SAN, as shown in Figure 3-19.

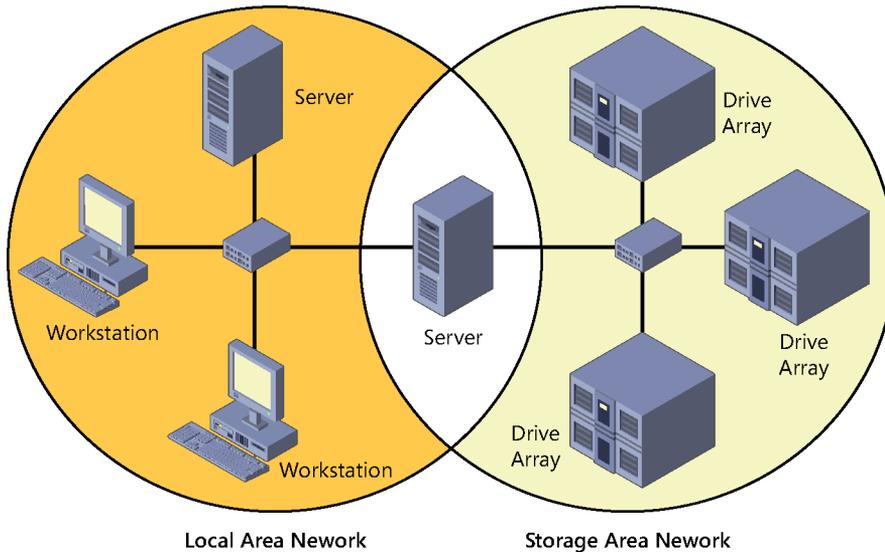


FIGURE 3-19 A server connected to a SAN.

The advantages of SANs are many. By connecting the storage devices to a network instead of to the servers themselves, you avoid the limitations imposed by the maximum number of devices you can connect directly to a computer. SANs also provide added flexibility in their communications capabilities. Because any device on a SAN can communicate with any other device on the same SAN, high-speed data transfers can occur in any of the following ways:

- **Server to storage** Servers can access storage devices over the SAN just as if they were connected directly to the computer.
- **Server to server** Servers can use the SAN to communicate directly with one another at high speeds to avoid flooding the LAN with traffic.
- **Storage to storage** Storage devices can communicate among themselves without server intervention, for example, to perform backups from one medium to another or to mirror drives on different arrays.

Although a SAN is not in itself a high-availability technology, you can make it one by connecting redundant servers to the same network, as shown in Figure 3-20, enabling them to access the same data storage devices. If one server should fail, another can assume its roles by accessing the same data. This is called server clustering.

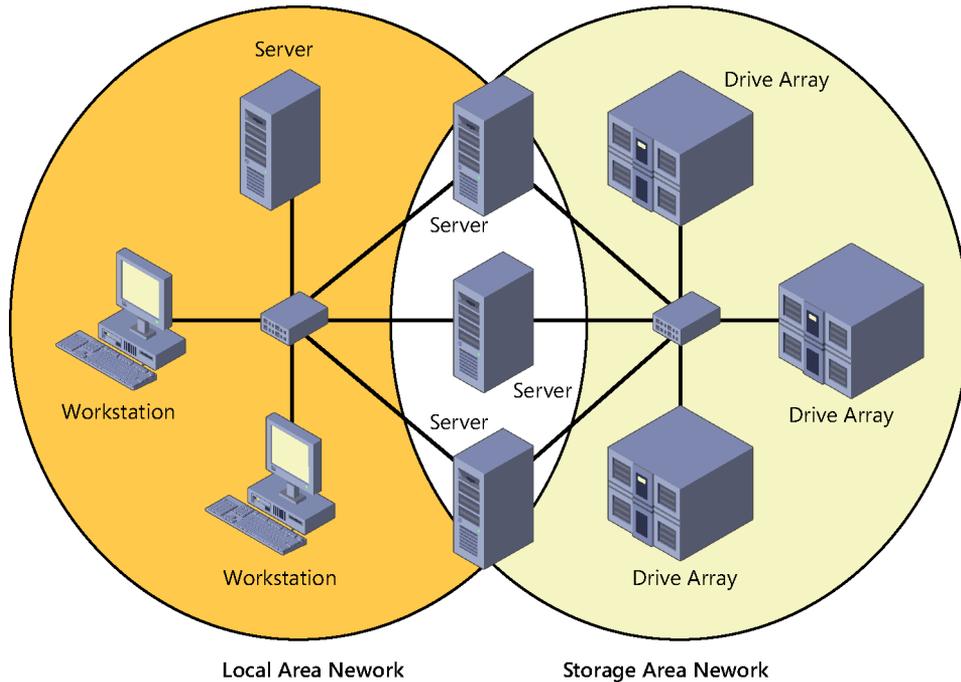


FIGURE 3-20 Multiple servers connected to a SAN.

Because they use standard networking technologies, SANs can also greatly extend the distances between servers and storage devices. You can design a SAN that spans different rooms, different floors, or even different buildings, just as you would a standard computer network.

Servers and storage devices cannot exchange SCSI commands over a SAN connection the way they do when the devices are directly connected using a SCSI cable. To communicate over a SAN, servers and storage devices map their SCSI communications onto another protocol, such as Fibre Channel.

Using Fibre Channel

Fibre Channel is a versatile SAN communications technology, supporting various network media, transmission speeds, topologies, and upper-level protocols. Its primary disadvantage is that it requires specialized hardware that can be extremely expensive.

MORE INFORMATION FIBRE CHANNEL

The nonstandard spelling of the word *fibre* in Fibre Channel is deliberate, to distinguish the term from fiber optic. Fibre Channel can run on either twisted-pair copper or optical cables, whereas the spelling *fiber* always refers to an optical medium.

Installing a Fibre Channel SAN entails building an entirely new network with its own special medium, switches, and network interface adapters. In addition to the hardware costs, which can easily be 10 times those of a traditional Ethernet network, there are also installation and maintenance expenses to consider. Fibre Channel is a rather esoteric technology, with relatively few experts in the field. To install and maintain a Fibre Channel SAN, an organization must either hire experienced staff or train existing personnel on the new technology.

Connecting virtual machines to a SAN

The specialized networking technologies used to build Fibre Channel SANs have, in the past, made it difficult to use them with virtualized servers. However, Windows Server 2012 Hyper-V now supports the creation of virtual Fibre Channel adapters.

A Hyper-V Fibre Channel adapter is essentially a pass-through device that enables a VM to access a physical Fibre Channel adapter installed in the computer, and through that, to access the external resources connected to the SAN. With this capability, applications running on VMs can access data files stored on SAN devices, and administrators can use VMs to create server clusters with shared storage subsystems.

To support virtual Fibre Channel connectivity, the physical Fibre Channel host bus adapter(s) in the host computer must have drivers that explicitly support virtual Fibre Channel. At the time of the release of Windows Server 2012, this support is relatively rare, but more manufacturers are expected to update their drivers to provide the necessary support. Your SAN must also be able to address its connected resources by using logical unit numbers (LUNs).

Assuming you have the appropriate hardware and software installed on the host computer, you implement the Fibre Channel capabilities in Hyper-V by first creating a virtual SAN by using the Virtual SAN Manager, accessible from Hyper-V Manager. When you create the virtual SAN, the World Wide Node Names (WWNNs) and World Wide Port Names (WWPNs) of your host bus adapter appear, as shown in Figure 3-21.

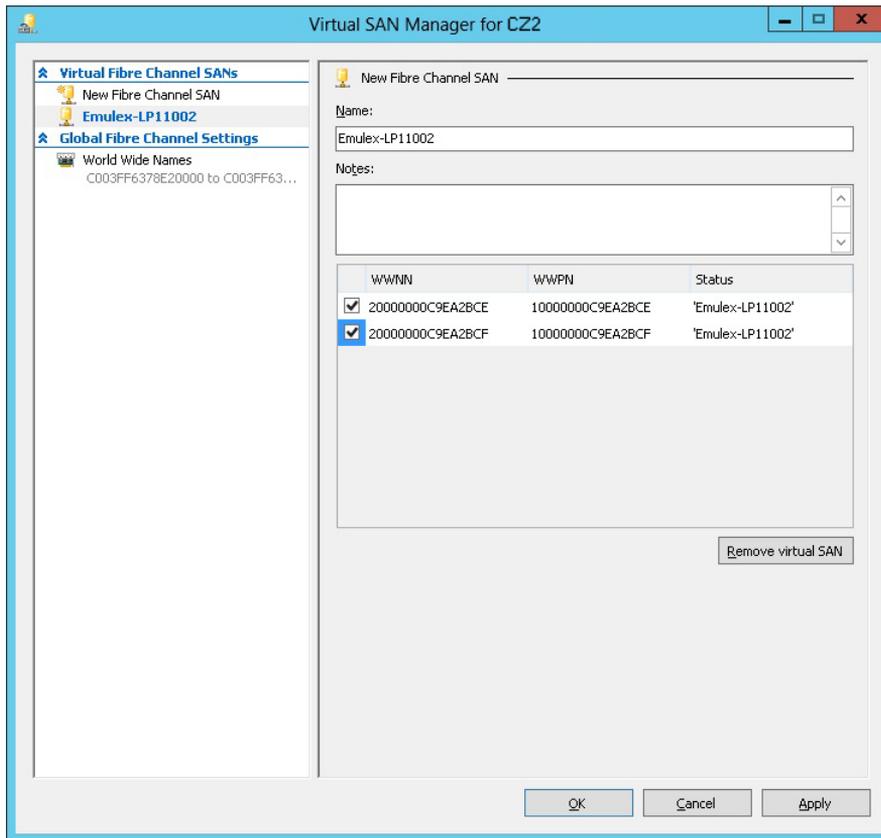


FIGURE 3-21 WWNNs and WWPNS in a virtual SAN.

The next step is to add a Fibre Channel adapter to a VM from the Add Hardware page in the Settings dialog box. When you do this, the virtual SAN you created earlier is available on the Fibre Channel Adapter page, shown in Figure 3-22. Hyper-V virtualizes the SAN and makes the WWNNs and WWPNS available to the VM.

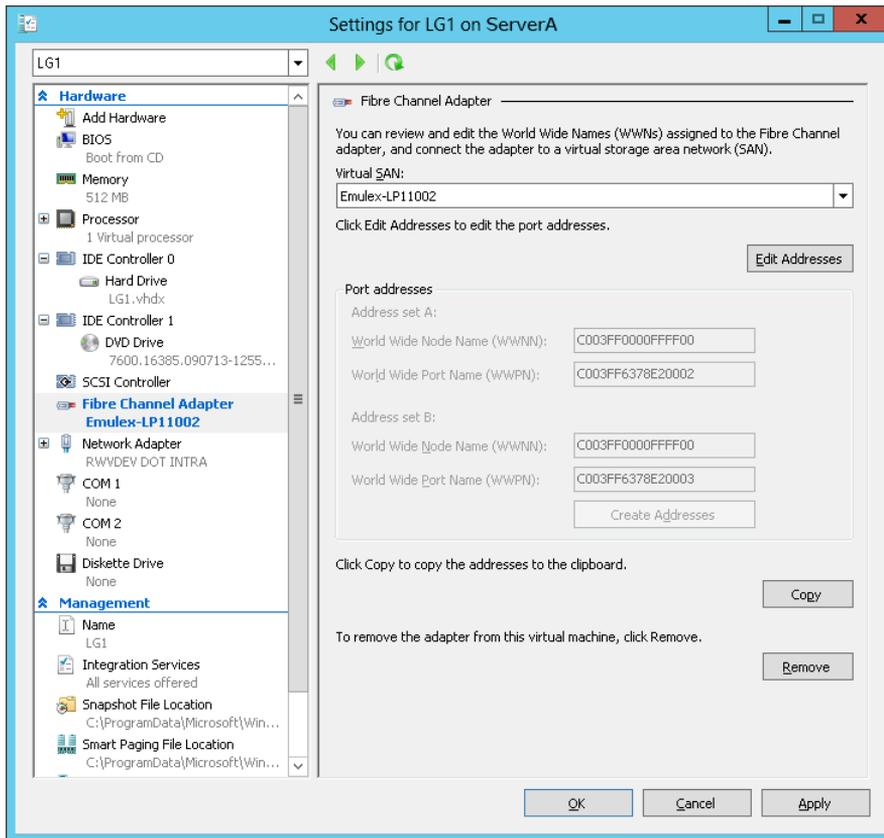


FIGURE 3-22 A Fibre Channel adapter in a VM.

Objective summary

- Hyper-V uses a specialized VHD format to package part of the space on a physical disk and make it appear to the VM as though it is a physical hard disk drive.
- A dynamic hard disk image is an image file with a specified maximum size, which starts small and expands as needed to accommodate the data the system writes to it.
- A differencing hard disk image is a child image file associated with a specific parent image. The system writes all changes made to the data on the parent image file to the child image, to facilitate a rollback at a later time.
- VHDX image files in Windows Server 2012 can be as large as 64 TB, and they also support 4-KB logical sector sizes to provide compatibility with new 4-KB native drives.
- A pass-through disk is a type of virtual disk that points not to an area of space on a physical disk, but to a physical disk drive installed on the host computer.

- In Hyper-V, a snapshot is a captured image of the state, data, and hardware configuration of a VM at a particular moment in time.
- The specialized networking technologies used to build Fibre Channel SANs have, in the past, made it difficult to use them with virtualized servers. However, Windows Server 2012 Hyper-V now supports the creation of virtual Fibre Channel adapters.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. Which of the following statements about VHDX files is not true?
 - A. VHDX files can be as large as 64 TB.
 - B. VHDX files can only be opened by computers running Windows Server 2012.
 - C. VHDX files support larger block sizes than VHD files.
 - D. VHDX files support 4-KB logical sectors.
2. Which of the following must be true about a pass-through disk?
 - A. A pass-through disk must be offline in the guest OS that will access it.
 - B. A pass-through disk must be offline in the parent partition of the Hyper-V server.
 - C. A pass-through disk can only be connected to a SCSI controller.
 - D. A pass-through disk must be added to a VM with the Disk Management snap-in.
3. The Merge function only appears in the Edit Virtual Hard Disk Wizard under which of the following conditions?
 - A. When you select a VHDX file for editing
 - B. When you select two or more disks for editing
 - C. When you select a disk with free space available in it
 - D. When you select a differencing disk for editing
4. Which of the following are valid reasons not to take snapshots of VMs? (Choose all that apply.)
 - A. Snapshots can consume a large amount of disk space.
 - B. Each snapshot requires a separate copy of the VM’s memory allocation.
 - C. Each snapshot can take several hours to create.
 - D. The existence of snapshots slows down VM performance.
5. Which of the following is not required to add a Fibre Channel adapter to a Hyper-V VM?
 - A. You must create a Fibre Channel virtual SAN.
 - B. You must have a physical Fibre Channel adapter installed in the host computer.

- C. You must have a Fibre Channel adapter driver that supports virtual networking.
- D. You must have a SCSI cable connecting the Fibre Channel adapter to the storage devices.



Thought experiment

In the following thought experiment, apply what you've learned about this objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

Ed wants to create a new VHD file on his Hyper-V server by using Windows PowerShell. He runs the `Get-Disk` cmdlet and receives the following results:

Number	Friendly Name	Operational Status	Total Size
0	WDC WD5003ABYX-18WERA0	Online	465.76 GB MBR
1	WDC WD1002FAEX-00Z3A0	Online	931.51 GB GPT

What command should Ed use to create a new 500-GB fixed VHD for his Server A VM, in the Windows Server 2012 format, using space from the 931-GB drive on his computer, and a 4,096-byte sector size?

Objective 3.3: Create and configure virtual networks

Networking is a critical part of creating a VM infrastructure. Depending on your network plan, the VMs you create on a Windows Server 2012 Hyper-V server can require communication with other VMs, with the computers on your physical network, and with the Internet.

When you build a network out of physical computers, you install a network interface adapter in each one and connect it to a hardware switch. The same principle is true in a Hyper-V environment, except that you use virtual components instead of physical ones. Each VM you create has at least one virtual network adapter, and you can connect that adapter to a virtual switch. This enables you to connect the VMs on your Hyper-V server in various network configurations that either include or exclude the systems on your physical network.

You can create multiple virtual switches on a Hyper-V server and multiple network adapters in each VM. This enables you to create a flexible networking environment that is suitable for anything from a laboratory or classroom network to a production environment. In addition, Windows Server 2012 has added the ability to create extensions for virtual switches so that software developers can enhance their capabilities.

This objective covers how to:

- Implement Hyper-V Network Virtualization
- Configure Hyper-V virtual switches
- Optimize network performance
- Configure MAC addresses
- Configure network isolation
- Configure synthetic and legacy virtual network adapters

Creating virtual switches

A virtual switch, like its physical counterpart, is a device that functions at Layer 2 of the Open Systems Interconnect (OSI) reference model. A switch has a series of ports, each of which is connected to a computer's network interface adapter. Any computer connected to the switch can transmit data to any other computer connected to the same switch.

Unlike physical switches, the virtual switches created by Hyper-V can have an unlimited number of ports, so administrators don't have to be concerned about connecting switches or about uplinks and crossover circuits.

Creating the default virtual switch

The Windows Server 2012 Add Roles and Features Wizard provides the opportunity to create virtual switches when you install the Hyper-V role. When you install Hyper-V on a server running Windows Server 2012, the Create Virtual Switches page provides you with the opportunity to create a virtual switch for each of the physical network adapters installed in the host computer. These switches enable VMs to participate on the networks to which the physical adapters are connected.

When you create a virtual switch in this manner, the networking configuration in the host OS on the parent partition changes. The new virtual switch appears in the Network Connections window, and if you examine its properties, you can see that the switch is bound to the operating system's TCP/IP client, as shown in Figure 3-23.

Meanwhile, Hyper-V also changes the properties of original network connection representing the physical network interface adapter in the computer. The physical network adapter is now bound only to the virtual switch, as shown in Figure 3-24.

As a result, the computer's physical network configuration, in which its network adapter is connected to an external physical switch, is overlaid by the virtual network configuration created by Hyper-V. In this virtual configuration, the virtual switch is connected to the physical switch and the network adapter in the host OS is connected to the virtual switch. The internal virtual network and the external physical network are joined into a single LAN, just as if you connected two physical switches.

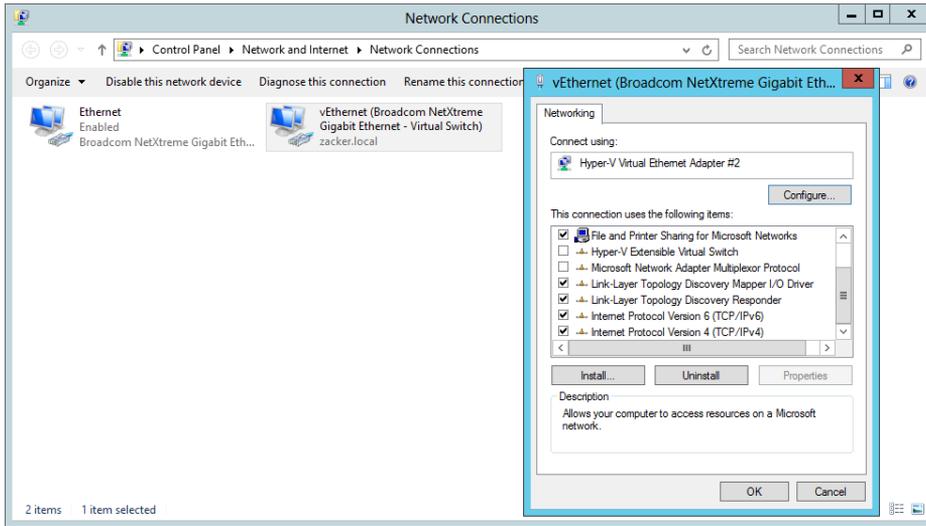


FIGURE 3-23 A virtual switch and its properties, displayed in the host OS.

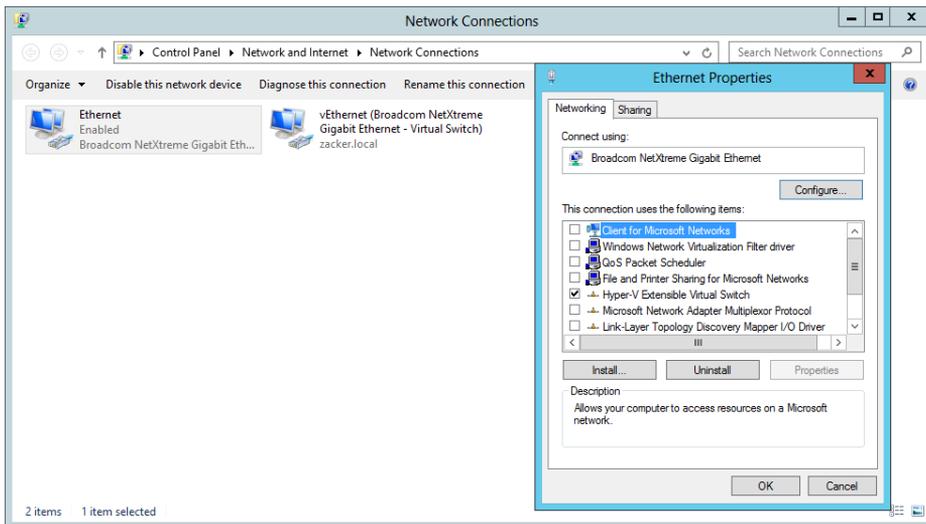


FIGURE 3-24 A network interface adapter in the host OS, bound to a virtual switch.

Once Hyper-V has created the virtual switch and made these configuration changes, any new VMs that administrators choose to connect to the virtual switch become part of this conjoined network, as do any physical computers connected to the physical network through an external switch.

This type of virtual switch is, in Hyper-V terminology, an external network switch because it provides connections external to the Hyper-V environment. This is typically the preferred

arrangement for a production network in which Hyper-V VMs provide and consume services for the entire network.

For example, a VM connected to this switch will automatically obtain an IP address from a Dynamic Host Configuration Protocol (DHCP) server on the physical network, if there is one. As an alternative, you could configure a VM as a DHCP server and let it provide addresses to all of the systems on the network, virtual or physical.

Perhaps more important, this arrangement can also enable your VMs to access the Internet by using the router and DNS servers on the external network. The VMs can then download OS updates from the Windows Update servers on the Internet, just as external machines often do.

There are situations in which this type of virtual switch is inappropriate. If you are creating a laboratory network for product testing or a classroom network, you might not want it to be accessible to or from the external network. In these cases, you must create a different type of virtual switch by using the Virtual Switch Manager in Hyper-V Manager.

Creating a new virtual switch

Hyper-V in Windows Server 2012 supports three types of switches, which you must create in the virtual Switch Manager before you can connect VMs to them.

To create a new virtual switch, use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager to open the Hyper-V Manager console.
3. In the left pane, select a Hyper-V server.
4. From the Actions pane, select Virtual Switch Manager. The Virtual Switch Manager dialog box for the Hyper-V server opens, as shown in Figure 3-25.

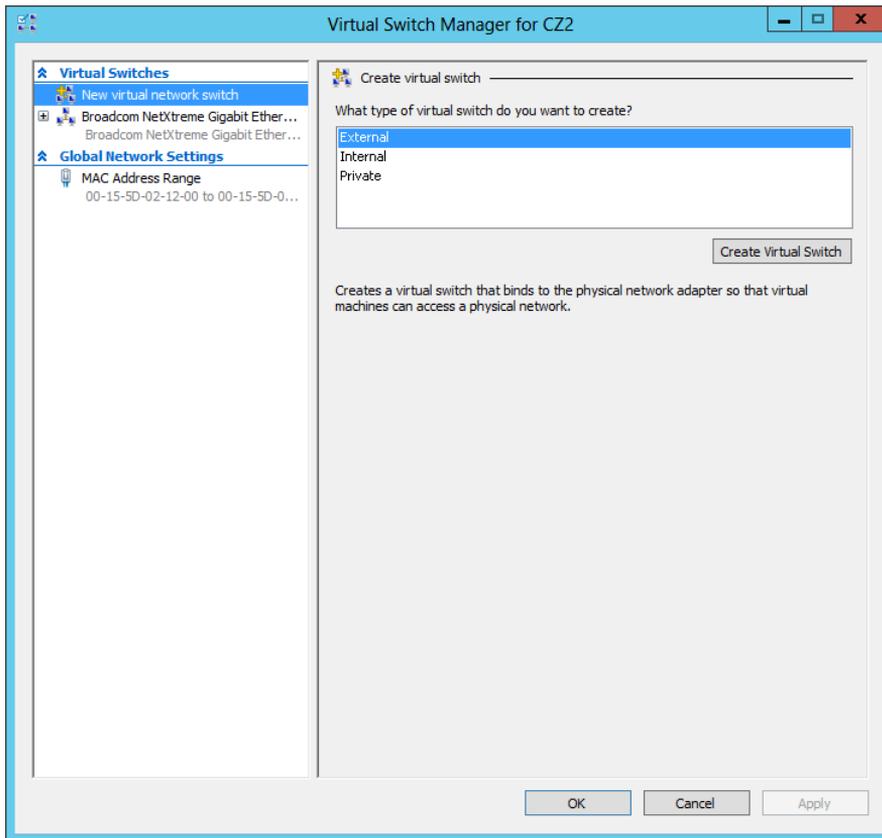


FIGURE 3-25 The Virtual Switch Manager dialog box.

5. In the Create Virtual Switch section, select one of the following switch types:
 - **External** The virtual switch is bound to the networking protocol stack in the host OS and connected to a physical network interface adapter in the Hyper-V server. VMs running on the server's parent and child partitions can all access the physical network to which the physical adapter is connected.
 - **Internal** An internal network switch is bound to a separate instance of the networking protocol stack in the host OS, independent from the physical network interface adapter and its connected network. VMs running on the server's parent and child partitions can all access the virtual network implemented by the virtual switch, and the host OS on the parent partition can access the physical network through the physical network interface adapter, but the VMs on the child partitions cannot access the physical network through the physical adapter.
 - **Private** A private network switch exists only in the Hyper-V server and is accessible only to the VMs running on the child partitions. The host OS on the parent

partition can access the physical network through the physical network interface adapter, but it cannot access the virtual network created by the virtual switch.

6. Click **Create Virtual Switch** to open the **Virtual Switch Properties** page.
7. Configure the following options, if desired:
 - **Allow Management Operating System To Share This Network Adapter** Selected by default when you create an external virtual switch, clearing this check box excludes the host OS from the physical network while allowing access to the child VMs.
 - **Enable Single Root I/O Virtualization (SR-IOV)** Enables you to create an external virtual switch that is associated with a physical network adapter capable of supporting SR-IOV. This option is only available when creating a new virtual switch; you cannot modify an existing virtual switch to use this option.
 - **Enable Virtual LAN Identification For Management Operating System** If your host computer is connected to a physical switching infrastructure that uses virtual LANs (VLANs) to create separate subnets, you can select this check box and enter a VLAN identifier to associate the virtual switch with a particular VLAN on your physical network.
8. Click **OK**. The new virtual switch appears in the left pane, in the list of virtual switches.

You can create additional virtual switches as needed. You can only create one external switch for each physical network adapter in the computer, but you can create multiple internal or private switches to create as many virtual networks as you need.

NOTE USING WINDOWS POWERSHELL

To create a new virtual switch by using Windows PowerShell, you use the `New-VMSwitch` cmdlet with the following basic syntax:

```
New-VMSwitch <switch name> -NetAdapterName <adapter name>  
[-SwitchType Internal|Private]
```

For example, to create an external switch called LAN Switch, you would use the following command:

```
New-VMSwitch "LAN Switch" -NetAdapterName "Ethernet"
```

Configuring MAC addresses

Every network interface adapter has a Media Access Control (MAC) address—sometimes called a hardware address—that uniquely identifies the device on the network. On physical network adapters, the MAC is assigned by the manufacturer and permanently entered in the adapter's firmware. The MAC address is a 6-byte hexadecimal value, the first three bytes of which are an organizationally unique identifier (OUI) that specifies the manufacturer, and the last three bytes of which identify the adapter itself.

The MAC address is essential to the operation of a LAN, so the virtual network adapters on a Hyper-V server need to have them. The server has at least one real MAC address, provided in its physical network adapter, but Hyper-V cannot use that one address for all the virtual adapters connecting VMs to the network.

Instead, Hyper-V creates a pool of MAC addresses during the installation of the role, and it assigns addresses from this pool to VMs as you create them. To view or modify the MAC address pool for the Hyper-V server, you open the Virtual Switch Manager and, under Global Network Settings, select MAC Address Range, as shown in Figure 3-26.

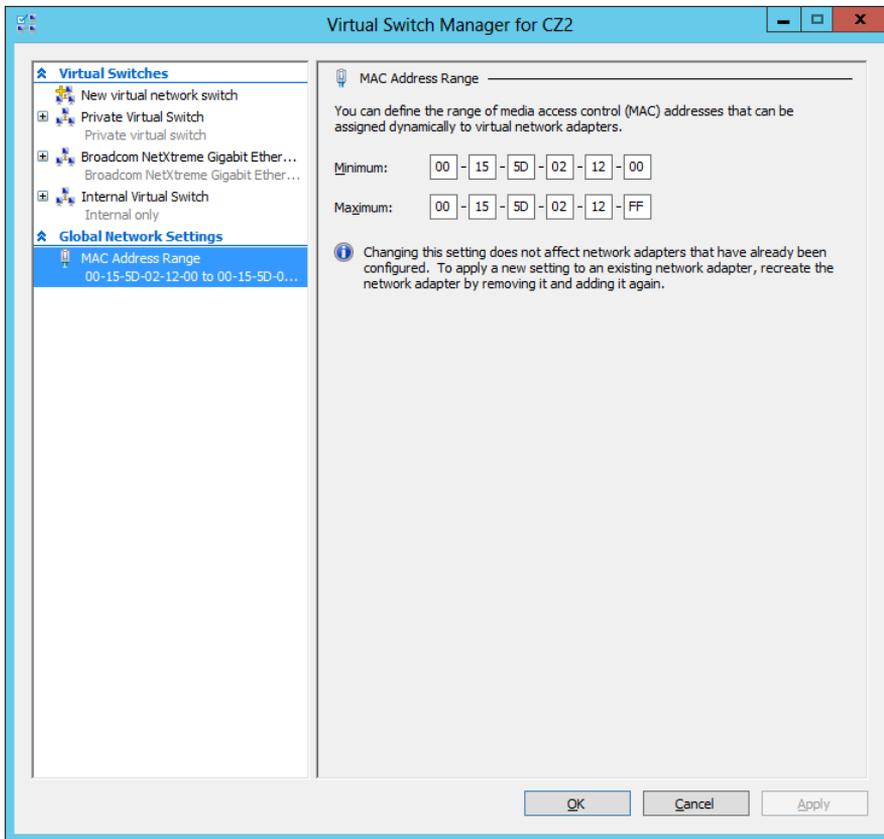


FIGURE 3-26 The MAC Address Range in the Virtual Switch Manager.

The first three bytes of the MAC address range are always 00-15-5D, which is an OUI registered by Microsoft. The fourth and fifth bytes of the MAC address are the last two bytes of the IP address assigned to the server's physical network adapter, converted to hexadecimals. The sixth and last byte of the MAC address contains the range of values from 00 to FF, which provides 256 possible addresses.

The Hyper-V server assigns the MAC addresses to the network adapters in VMs as administrators create the adapters. The adapters retain their MAC addresses permanently or until the adapter is removed from the VM. The server reclaims any unused addresses and reuses them.

The default pool of 256 addresses is expected to be sufficient for most Hyper-V VM configurations, but if it is not, you can modify the Minimum and Maximum values to enlarge the pool. To prevent address duplication, you should change the second-to-last byte only, making it into a range of addresses like the last byte.

For example, if the range illustrated in Figure 3-26 provides 256 addresses with the following values:

00-15-1D-02-12-00 to 00-15-1D-02-12-FF

Modifying only the least significant digit, as in the following values, will increase the pool from 256 to 4,096, as follows:

00-15-1D-02-10-00 to 00-15-1D-02-1F-FF

WARNING MAC ADDRESSES

When you modify the MAC address pool and you have other Hyper-V servers on your network, you must be careful not to create the chance for duplicate MAC addresses, or networking problems can occur.

Creating virtual network adapters

Once you have created virtual switches in Hyper-V Manager, you can connect VMs to them by creating and configuring virtual network adapters. When you create a new VM, the default configuration includes one virtual network adapter. The New Virtual Machine Wizard includes a Configure Networking page, on which you can select one of the virtual switches you have created.

If you have created only the default external virtual switch when installing Hyper-V, then connecting a VM to that switch joins the system to the physical network. If you want to create additional network adapters in your VMs, you must use the following procedure.

1. Log on to the server running Windows Server 2012 using an account with administrative privileges. The Server Manager window opens.
2. From the Tools menu, select Hyper-V Manager to open the Hyper-V Manager console.
3. In the left pane, select a Hyper-V server.
4. In the Virtual Machines list, select a VM and, in the Actions pane, click Settings. The Settings dialog box for the VM appears.
5. In the Add Hardware list, select Network Adapter and click Add. A new adapter appears in the Hardware list, as shown in Figure 3-27.

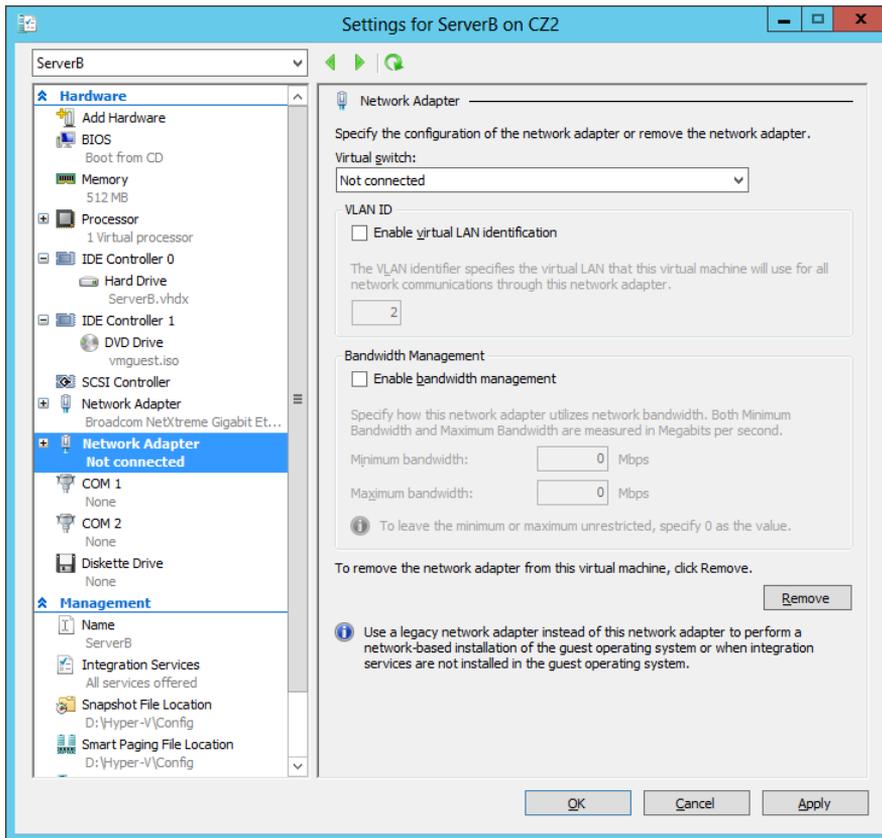


FIGURE 3-27 A new network adapter in the Settings dialog box.

6. In the Virtual Switch drop-down list, select the switch to which you want to connect the network adapter.
7. If your host computer is connected to a physical switching infrastructure that uses VLANs to create separate subnets, you can select the Enable Virtual LAN Identification check box and enter a VLAN identifier to associate the network adapter with a particular VLAN on your physical network.
8. To control the amount of network bandwidth allocated to the network adapter, select the Enable Bandwidth Management check box and supply values for the Minimum Bandwidth and Maximum Bandwidth settings.
9. Click OK. The settings are saved to the VM configuration.

You can create up to 12 network adapters on a Windows Server 2012 Hyper-V server: eight synthetic and four emulated.

Synthetic adapters and emulated adapters

Selecting the Network Adapter option on the Add Hardware page creates what is known in Hyper-V terminology as a synthetic network adapter. Hyper-V supports two types of network and storage adapters: synthetic and emulated (sometimes called legacy).

A synthetic adapter is a purely virtual device that does not correspond to a real-world product. Synthetic devices in a VM running on a child partition communicate with the parent partition by using a high-speed conduit called the VMBus.

The virtual switches you create in Hyper-V reside in the parent partition and are part of a component called the network Virtualization Service Provider (VSP). The synthetic network adapter in the child partition is a Virtualization Service Client (VSC). The VSP and the VSC are both connected to the VMBus, which provides interpartition communications, as shown in Figure 3-28. The VSP, in the parent partition, provides the VSC, in the child partition, with access to the physical hardware in the host computer; that is, the physical network interface adapter.

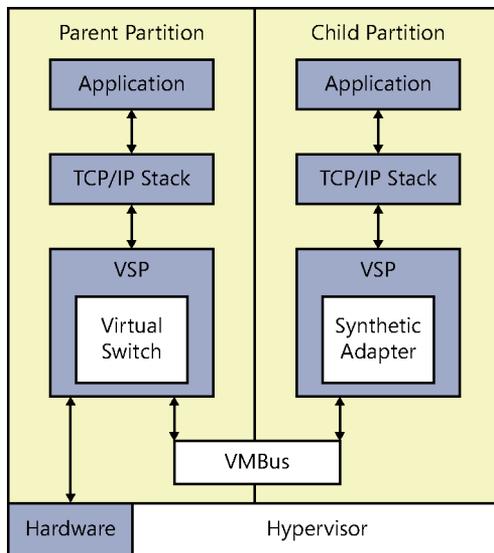


FIGURE 3-28 Synthetic network adapters communicate by using the VMBus.

Because they have access to the hardware through the VMBus, synthetic adapters provide a much higher level of performance than the alternative, emulated adapters. Synthetic adapters are implemented as part of the Guest Integration Services package that runs on supported guest OSs. The main drawback of synthetic network adapters is that they are not operational until the OS is loaded on the VM.

An emulated adapter—sometimes called a legacy adapter—is a standard network adapter driver that communicates with the parent partition by making calls directly to the hypervisor, which is external to the partitions, as shown in Figure 3-29. This communication method is

substantially slower than the VMBus used by the synthetic network adapters, and is therefore less desirable.

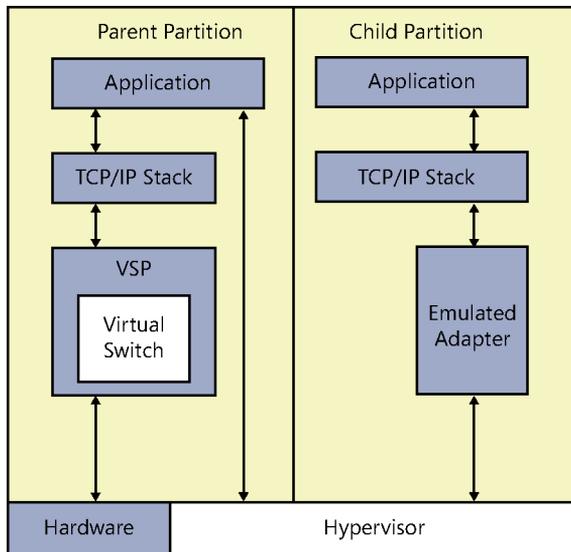


FIGURE 3-29 Emulated network adapters communicate by using the hypervisor.

To install an emulated adapter, you use the same procedure described earlier, except that you select Legacy Network Adapter from the Add Hardware list. Unlike synthetic adapters, emulated adapters load their drivers before the OS, so it is possible to boot the VM by using the Preboot eXecution Environment (PXE) and deploying an OS over the network.

This is one of two scenarios in which using an emulated adapter is preferable to using a synthetic adapter. The other is when you are installing an OS on your VMs that does not have a Guest Integration Services package available for it.

Configuring hardware acceleration settings

Some physical network interface adapters have features that are designed to improve performance by offloading certain functions from the system processor to components built into the adapter itself. Hyper-V includes support for some of these features, as long as the hardware in the physical network adapter supports them properly.

When you expand a network adapter in the Setting dialog box of a VM, you gain access to the Hardware Acceleration page. On this page, you can configure the following hardware acceleration settings:

- **Enable Virtual Machine Queue** Virtual machine queue (VMQ) is a technique that stores incoming packets intended for VMs in separate queues on the physical network adapter and delivers them directly to the VMs, bypassing the processing normally performed by the virtual switch on the parent partition.

- **Enable IPsec Task Offloading** Uses the components on the network adapter to perform some of the cryptographic functions required by IPsec. You can also specify the maximum number of security associations you want the adapter to be able to calculate.
- **Single-Root I/O Virtualization** Enables the virtual adapter to take advantage of the SR-IOV capabilities of the physical adapter.

Configuring advanced network adapter features

The Advanced Features page provides additional options for supporting network adapter capabilities, as follows:

- **Static MAC Address** By default, virtual network adapters receive a dynamically assigned MAC address from the Hyper-V server. However, you can opt to create a static MAC address by using this option. The only requirement is that no other adapter, virtual or physical, on the same network uses the same address.
- **Enable MAC Address Spoofing** When enabled, the port in the virtual switch to which the virtual network adapter is connected can send and receive packets that contain any MAC address. The virtual switch port can also learn of new MAC addresses and add them to its forwarding table.
- **Enable DHCP Guard** Prevents the adapter from processing messages sent by rogue DHCP servers.
- **Port Mirroring Mode** Enables the adapter to forward all the packets it receives over the network to another virtual adapter for analysis by using an application such as Network Monitor.
- **NIC Teaming** Enables the adapter to add its bandwidth to that of other adapters in the same guest OS in a NIC teaming arrangement.

Creating virtual network configurations

Hyper-V makes it possible to extend nearly any existing physical network configuration into its virtual space or create a completely separated and isolated network within the Hyper-V environment.

The basic default configuration of a Hyper-V VM connects its network adapter to an external virtual switch, thus attaching the guest OS on the VM to the outside network. The VM can then take advantage of services running on the outside network and send traffic through routers to other networks, including the Internet.

This type of arrangement can enable administrators to consolidate many physical servers into VMs on a single Hyper-V server, providing them all with access to the entire network. There is no distinction here between the physical network and the virtual one in the Hyper-V space.

Extending a production network into virtual space

Keep in mind that a Hyper-V server can have multiple physical network interface adapters installed in it, which might be connected to different networks to separate traffic or to the same network to increase available bandwidth. You might also have adapters dedicated to SAN connections for shared storage and server clustering.

Microsoft recommends the use of at least two physical network adapters in a Hyper-V server, with one adapter servicing the parent partition and the other connected to the child partitions. When you have more than two physical adapters in the server, you can create separate external virtual network switches for the physical adapters and connect each one to a separate VM.

Creating an isolated network

For testing and evaluation purposes or for classroom situations, administrators might want to create isolated network environments. By creating internal or private virtual switches, you can create a network that exists only within the Hyper-V space, with or without the parent partition included.

An isolated network such as this suffers from the weaknesses of its strengths. If you want to install the guest OSs by using Windows Deployment Services or configure the VMs by using DHCP, you must install and configure those services on your private network. The guest OSs also do not have access to the Internet, which prevents them from downloading OS updates. In this case, you must deploy appropriate substitutes on the private network.

One way to provide your systems with updates is to install two network adapters on each of your VMs, connecting one to a private switch and one to an external switch. This enables the VMs to access the Internet and the private network.

Another method for creating an isolated network is to use VLANs. This is particularly helpful if you have VMs on different Hyper-V servers that you want to add to the isolated network. By connecting the network adapters to an external switch and configuring them with the same VLAN identifier, you can create a network within a network, which isolates the VLAN from other computers. You can, for example, deploy a DHCP server on your VLAN without it interfering with the other DHCP servers in your production environment.

Objective summary

- Networking is a critical part of creating a VM infrastructure. Depending on your network plan, the VMs you create on a Windows Server 2012 Hyper-V server can require communication with other VMs, with the computers on your physical network, and with the Internet.
- A virtual switch, like its physical counterpart, is a device that functions at Layer 2 of the OSI reference model. A switch has a series of ports, each of which is connected to a computer's network interface adapter. Any computer connected to the switch can transmit data to any other computer connected to the same switch.

- Hyper-V in Windows Server 2012 supports three types of switches: external, internal, and private, which you must create in the virtual Switch Manager before you can connect VMs to them.
- Every network interface adapter has a MAC address—sometimes called a hardware address—that uniquely identifies the device on the network.
- Once you have created virtual switches in Hyper-V Manager, you can connect VMs to them by creating and configuring virtual network adapters.
- Selecting the Network Adapter option on the Add Hardware page creates what is known in Hyper-V terminology as a synthetic network adapter. Hyper-V supports two types of network and storage adapters: synthetic and emulated (sometimes called legacy).

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of this chapter.

1. Which of the following are valid reasons for using an emulated network adapter rather than a synthetic one? (Choose all that apply.)
 - A. You want to install the guest OS by using a Windows Deployment Services server.
 - B. There is no Guest Integration Services package available for the guest OS you plan to use.
 - C. The manufacturer of your physical network adapter has not yet provided a synthetic network adapter driver.
 - D. The emulated network adapter provides better performance.
2. Which of the following statements is not true about synthetic network adapters?
 - A. Synthetic adapters communicate with the parent partition by using the VMBus.
 - B. Synthetic adapters require the Guest Integration Services package to be installed on the guest OS.
 - C. Synthetic adapters provide faster performance than emulated adapters.
 - D. Synthetic adapters can start the child VM by using a PXE network boot.
3. What is the maximum number of ports supported by a Hyper-V virtual switch?
 - A. 8
 - B. 256
 - C. 4,096
 - D. Unlimited

4. Which of the following virtual switch types does not enable guest OSs to communicate with the parent partition?
 - A. External
 - B. Internal
 - C. Private
 - D. Isolated
5. How many dynamically assigned MAC addresses can a Hyper-V server provide by default?
 - A. 8
 - B. 256
 - C. 4,096
 - D. Unlimited



Thought experiment

In the following thought experiment, apply what you've learned about this objective to predict what steps you need to take. You can find answers to these questions in the "Answers" section at the end of this chapter.

Ralph has a Windows Server 2012 Hyper-V server with one physical network adapter and one external virtual switch connected to that adapter. This arrangement enables the VMs on the server to automatically download OS updates from the Internet. However, Ralph wants to use the VMs on the Hyper-V server to construct an isolated test network on which he can evaluate new software products. The test network must have its own DHCP server that does not interfere with the DHCP server on the production network.

How can Ralph construct the test network he needs from his VMs without changing the configuration that provides the machines with Internet access?

Answers

This section contains the solutions to the thought experiments and answers to the lesson review questions in this chapter.

Objective 3.1: Review

1. **Correct answers:** B, C
 - A. **Incorrect:** In Type I virtualization, the hypervisor does not run on top of a host OS.
 - B. **Correct:** A Type I hypervisor runs directly on the computer hardware.
 - C. **Correct:** A Type II hypervisor runs on top of a host OS.
 - D. **Incorrect:** In Type II virtualization, the hypervisor does not run directly on the computer hardware.
2. **Correct answer:** A
 - A. **Correct:** Type I virtualization provides the best performance because the hypervisor runs directly on the computer hardware and does not have the overhead of a host OS.
 - B. **Incorrect:** Type II virtualization provides poorer performance than Type I because of the need to share processor time with the host OS.
 - C. **Incorrect:** Presentation virtualization is the term used to describe the Remote Desktop Services functionality in Windows. It is not designed for virtualizing servers.
 - D. **Incorrect:** RemoteApp is a technology for virtualizing individual applications and deploying them by using Remote Desktop Services.
3. **Correct answer:** B
 - A. **Incorrect:** Hyper-V Server does not include a license for any virtual instances.
 - B. **Correct:** Windows Server 2012 Datacenter edition includes a license that enables you to create an unlimited number of virtual instances.
 - C. **Incorrect:** Windows Server 2012 Standard edition includes a license that enables you to create two virtual instances.
 - D. **Incorrect:** Windows Server 2012 Foundation edition does not include support for Hyper-V.
4. **Correct answers:** A, B
 - A. **Correct:** Smart paging enables a VM to restart even if the amount of RAM specified as the startup value is unavailable. Smart paging causes the system to use disk space as a temporary substitute for memory during a system restart.
 - B. **Correct:** Dynamic Memory enables you to specify a minimum RAM value that is smaller than the startup RAM value, but Smart paging enables the system to function with those parameters.

- C. Incorrect:** Windows Memory Weight controls the allocation of memory among VMs, but it does not affect the ability of a system to start.
 - D. Incorrect:** Guest Integration Services is the software package that runs on a guest OS and facilitates communication with the parent partition.
- 5. Correct answer: C**
- A. Incorrect:** The instance of the OS on which you install Hyper-V does not become the hypervisor.
 - B. Incorrect:** The instance of the OS on which you install Hyper-V does not become the VMM.
 - C. Correct:** The instance of the OS on which you install the Hyper-V role becomes the parent partition.
 - D. Incorrect:** The instance of the OS on which you install the Hyper-V role does not become the child partition.

Objective 3.1: Thought experiment

Alice can enable Dynamic Memory on each of the eight VMs and set the minimum RAM value on each to 512 MB. This will enable each VM to start with 1,024 MB of memory and then reduce its footprint, allowing the next machine to start.

Objective 3.2: Review

- 1. Correct answer: B**
- A. Incorrect:** VHDX files can be as large as 64 TB, whereas VHD files are limited to 2 TB.
 - B. Correct:** Windows Server 2012 and Windows 8 can both open VHDX files.
 - C. Incorrect:** VHDX files support block sizes as large as 256 MB.
 - D. Incorrect:** VHDX files can support the 4,096-byte block sizes found on some newer drives.
- 2. Correct answer: B**
- A. Incorrect:** A pass-through disk must be online in the guest OS that will access it.
 - B. Correct:** A pass-through disk must be offline in the parent container so that the guest OS can have exclusive access to it.
 - C. Incorrect:** A pass-through disk can be connected to any type of controller.
 - D. Incorrect:** You do not use the Disk Management snap-in to add a pass-through disk to a VM; you use Hyper-V Manager.
- 3. Correct answer: D**
- A. Incorrect:** You can merge VHD or VHDX disks.
 - B. Incorrect:** You can only select one disk for editing.

- C. Incorrect:** There is no free space requirement when merging a disk.
 - D. Correct:** The Merge function appears only when you select a differencing disk for editing. The object of the function is to combine the data in the differencing disk with that of the parent.
- 4. Correct answers:** A, D
- A. Correct:** Snapshots consume disk space that could be better used for other purposes.
 - B. Incorrect:** Snapshots do not require a duplicate memory allocation.
 - C. Incorrect:** Under typical conditions, snapshots do not take several hours to create.
 - D. Correct:** The Hyper-V server must locate and process snapshots each time it accesses a VM's disk drives, slowing down its performance.
- 5. Correct answer:** D
- A. Incorrect:** You must create a Fibre Channel SAN before you can add a Fibre Channel adapter to a VM.
 - B. Incorrect:** You must have a physical Fibre Channel adapter before you can create virtual Fibre Channel components.
 - C. Incorrect:** The driver for your physical Fibre Channel adapter must support virtual networking.
 - D. Correct:** SCSI cables are not required for Fibre Channel installations.

Objective 3.2: Thought experiment

Ed should use the following Windows PowerShell command to create the VHD.

```
New-VHD -Path c:\servera.vhdx -Fixed -SizeBytes 500GB -LogicalSectorSizeBytes 4096 -SourceDisk 1
```

Objective 3.3: Review

- 1. Correct answers:** A, B
- A. Correct:** A Windows Deployment Server installation requires the network adapter to support PXE, which emulated adapters do, but synthetic adapters do not.
 - B. Correct:** Synthetic adapter drivers are installed as part of the Guest Integration Services package; if there is no package for the guest OS, then there are no synthetic drivers.
 - C. Incorrect:** Synthetic adapter drivers are not provided by hardware manufacturers.
 - D. Incorrect:** Synthetic adapters provide better performance than emulated adapters.

2. **Correct answer:** D
- A. **Incorrect:** Synthetic adapters use the faster VMBus for communications with the parent partition; emulated adapters must use calls to the hypervisor.
 - B. **Incorrect:** Synthetic adapter drivers are installed as part of the Guest Integration Services package on the guest OS.
 - C. **Incorrect:** Because of their more efficient communication with the parent partition, synthetic adapters perform better than emulated adapters.
 - D. **Correct:** Synthetic network adapters load with the Guest Integration Services on the guest OS, which prevents them from supporting PXE.
3. **Correct answer:** D
- A. **Incorrect:** Switches limited to eight connections would be insufficient for many Hyper-V installations.
 - B. **Incorrect:** Hyper-V switches are not limited to 256 connections.
 - C. **Incorrect:** Hyper-V switches are not limited to 4,096 connections.
 - D. **Correct:** Hyper-V virtual switches can support an unlimited number of connections.
4. **Correct answer:** C
- A. **Incorrect:** External switches enable the guest OSs to communicate with the outside network and the parent partition.
 - B. **Incorrect:** Internal switches enable the guest OSs to communicate with the parent partition but not with the outside network.
 - C. **Correct:** Private switches enable the guest OSs to communicate with one another but not with the outside network or the parent partition.
 - D. **Incorrect:** Isolated is not a technical term referring to a type of virtual switch.
5. **Correct answer:** B
- A. **Incorrect:** A pool of eight MAC addresses would be insufficient for many Hyper-V installations.
 - B. **Correct:** A Hyper-V server provides a pool of 256 MAC addresses by default. You can create more by modifying the default address range.
 - C. **Incorrect:** Hyper-V, by default, dedicates only one byte of the MAC address to a dynamic value, which is not enough to support 4,096 addresses.
 - D. **Incorrect:** Hyper-V creates a finite pool of MAC addresses by specifying minimum and maximum address values.

Objective 3.3: Thought experiment

Ralph can create an isolated test environment without changing the virtual switch configuration by selecting the Enable Virtual LAN Identification check box on the network adapter in each VM and specifying the same VLAN identifier for each VM he wants on the test network.