

THE MICROSOFT JOURNAL FOR IT PROFESSIONALS

TechNet

MAGAZINE

2



September 2013

- 2 **Editor's note**
Secure new platforms
Lafe Low

- 4 **Windows 7**
Manage disks and file systems
Jorge Orchilles

- 10 **Active Directory**
Protect your Active Directory data
Darren Mar-Elia

- 15 **Security**
Don't forget the basics
John Vacca

- 19 **Cloud Computing**
Challenges in cloud configuration
Dan Marinescu

- 22 **Windows 8**
File History explained
Bohdan Raciborski

- 26 **Identity Management**
Use two-factor authentication to mitigate fraud
Tom Jones, and Dan Griffin

- 30 **Windows 8**
Painless printing
Adrian Lannin

- 35 **SharePoint**
Secure SharePoint content
Dan Sullivan

- 39 **Windows Confidential**
Respect the brand
Raymond Chen

Editor's Note

Secure new platforms

Besides contending with an evolving cast of new viruses, malware, and other threats, you have to secure new platforms and new technologies.

Lafe Low

The security “threatscape” is constantly changing and evolving, with new viruses, new malware, and new threat vectors hackers can follow in their incessant attempts to compromise your systems. As if that wasn't enough, sweeping infrastructure changes have altered how you have to approach security. The old security techniques and tactics won't work as well on newer technologies such as mobile devices, cloud computing and virtual infrastructure.

Securing the cloud is a challenge simply because of where data is stored, how often it moves and who will have access to it. You have to ensure data is secured at all stages—when it's in use, in transit and at rest. Choose your provider carefully. Choose one who will work with you to ensure the level of security you need in order to sleep at night, and one who will share the liability should something go wrong.

The openness of cloud infrastructures and collaboration platforms such as SharePoint makes them attractive targets for data thieves. “To understand the risks to your content, it helps to think like an attacker,” says Dan Sullivan in this month's feature, “Secure SharePoint content.” He goes on to describe several attack vectors, including browser cache mining, endpoint malware, and careless or malicious employees.

Securing mobile devices brings up a whole handful of new and expanded security concerns. You have to grapple with remote ID management to ensure users logging into the corporate network are indeed who they claim to be; apply corporate policies to bring your own device (BYOD) devices; and maintain a schedule of OS and antivirus updates. There will be certain occasions when you simply can't apply corporate policy to a remote device that doesn't belong to the company. Remote wipe is a perfect example. If your star salesman loses his gilded smartphone, you can't apply a remote wipe. What about all his personal contacts and personal data?

Securing virtual systems brings up another new model for ensuring security. Certainly the process of applying new policies and patches is simplified. You can have them applied as needed whenever a registered user fires up his system at the beginning of the day. You can also start off each day with a clean system, thereby ensuring no virus or malware can linger.

So certain aspects of securing the expanding universe of devices have gotten easier, while other aspects have become a profound challenge. Keeping one step ahead of the changing security landscape is the best way to avoid those sleepless nights and late-night texts and phone calls.

Sleep well

Do you sleep well as an IT guy these days? If you do, that's probably the best sign that you have a firm grip on your organization's infrastructure. If you're tossing and turning, what is it that's keeping you up at night? What are you missing? What's your current focus?

Is there anything we can do with *TechNet Magazine* to help you get a good night's sleep? Let us know what you need, and we'll help you get the most out of your Microsoft technology investments. Sign up for our [LinkedIn group](#), send us an e-mail at tnmag@microsoft.com or e-mail me directly.



Lafe Low is the editor in chief of *TechNet Magazine*. A veteran technology journalist, he's also the former executive editor of 1105 Media's Redmond magazine.

Windows 7

Manage disks and file systems

You can configure and reconfigure disks, files, partitions and volumes to suit your specific needs.

Jorge Orchilles

Adapted from "Microsoft Windows 7 Administrator's Reference" (Syngress, an imprint of Elsevier).

The disks and file system are at the center of Windows 7 operations. It's important to have them correctly configured, or else your system might not function properly or at all. The disks are where all your information is stored: OS files, applications, data, everything. The file system determines how this information is stored on your disks.

Partitions

Partitions divide and segment your disks. You can have a disk with either a single partition or multiple partitions. Even though having a single partition is the simplest way to configure your disks, there are several reasons you might want to have more than one partition.

Having multiple partitions helps you separate your OS files, application files and data files. Sometimes, you'll need multiple partitions because of size limitations for a partition. You may also need multiple partitions to run a multi-boot system. This would especially be the case if different OSes used different file systems. You'd need a different partition for each of the different file systems.

MBR and GPT

Most legacy disks are Master Boot Record (MBR) disks. MBR disks store partition information in the MBR, hence the name. This information is generally stored in the first sector of the disk.

GUID Partition Table (GPT) disks store partition information in the GPT header. For compatibility with MBR systems, GPT disks continue to store the MBR entry as the first sector of the disk. Following this entry is the start of the GPT, also called the primary partition table header. For redundancy, the GPT header and partition table are also written at the end of the disk.

In Windows 7, you can have MBR or GPT disks. When you first add a new disk, you must choose one or the other. It's important you understand the differences between these two types of disks. Because the GPT format is newer, you may run into compatibility issues if you choose this format. MBR disks have a wider range of compatibility. GPT disks, however, support larger partition sizes.

Convert from MBR to GPT

You can convert an MBR disk to a GPT disk. To do so, the disk can't have any volumes. If the disk has volumes, you should remove them before starting the conversion process. To convert an MBR disk to a GPT disk, follow these steps:

1. Open the Disk Management Microsoft Management Console (MMC) snap-in.
2. Right-click on the disk and select Convert to GPT disk.
3. The disk will be converted, and should show as Online.

You can also convert a GPT disk to an MBR disk by following these steps:

1. Open the Disk Management MMC snap-in.
2. Right-click on the disk and select Convert to MBR disk.
3. The disk will be converted, and should show as Online.

Basic and dynamic disks

There are two types of disks available in Windows 7: basic disks and dynamic disks. You can think of basic disks as the traditional technology used for Windows disks. When disks are first created, they're created as basic disks. In the original Disk Creation wizard, you can convert the disk to a dynamic disk. You can also convert to a dynamic disk later.

Follow these steps to convert a basic disk to a dynamic disk:

1. Open the Disk Management MMC snap-in.
2. Right-click on the disk and select Convert to Dynamic Disk. The Convert to Dynamic Disk window will appear.
3. Select the basic disks you want to convert to dynamic disks.
4. Click OK. The disks are converted to dynamic disks.

Converting between MBR and GPT format

When you convert a disk from MBR to GPT format, you'll notice the amount of unallocated disk space on the disk decreases. This is because of the additional space used to hold disk and partition information on GPT disks. When you convert from GPT format to MBR format, you'll notice the reverse effect. If the disk has a volume on it, you can't convert it back to a basic disk. You must delete the volume before attempting to convert the disk back to a basic disk.

To convert a dynamic disk back to a basic disk, follow these steps:

1. Open the Disk Management MMC snap-in.
2. Right-click on the disk and select Convert to Basic Disk.
3. Click OK. The disk is converted back to a basic disk.

Volumes

Once you have your disks created and configured, you need to create volumes. You can't store any information on your disks until you've created volumes. There are several different types of volumes from which to choose. You can create simple volumes, spanned volumes, striped volumes or mirrored volumes. Creating a simple volume is easy. Just follow these steps:

1. Right-click on "Unallocated space." Select New Simple Volume.
2. The New Simple Volume wizard appears. Click Next.
3. On the Specify Volume Size screen, enter the desired size of the new volume. Click Next.
4. On the Assign Drive Letter or Path screen, choose the desired drive letter for the new volume. Click Next.
5. On the Format Partition screen, choose the option for "Format this volume using the following settings."
6. Choose the NTFS file system. Enter a volume label. Select the option for Perform a Quick Format. Click Next.
7. On the Completing the new Simple Volume Wizard screen, click Finish. The new volume will be formatted and should show a status of Healthy.

To create a new spanned volume, follow these steps:

1. Right-click on "Unallocated space." Select New Spanned Volume.
2. The New Spanned Volume wizard appears. Click Next.
3. The Select Disks screen will appear. Select the disk you want to add to the spanned volume. Click Add.
4. Once the disk is added, you can specify how much disk space from the disk you want to add to the spanned volume. Click Next.
5. On the Assign Drive Letter or Path screen, choose the desired drive letter for the new volume. Click Next.
6. On the Format Partition screen, choose the option for "Format this volume using the following settings."
7. Choose the NTFS file system. Enter a volume label. Select the option for Perform a Quick Format. Click Next.
8. On the Completing the new Simple Volume Wizard screen, click Finish. The new volume will be formatted and should show a status of Healthy on each of the disks to which it was added.

Striped volumes and spanned volumes are similar, but have one big difference: Both types of volumes can stretch across multiple disks. Striped volumes use the same amount of disk space on all disks, whereas spanned volumes can use a different amount of space on each disk. To create a new striped volume, follow these steps:

1. Open the Disk Management MMC snap-in.
2. Right-click "Unallocated space." Select New Striped Volume.
3. The New Striped Volume wizard appears. Click Next.
4. The Select Disks screen appears. Select the disks you want to add to the spanned volume. Click Add.
5. Once the disks are added, you can specify how much disk space from the disk you want added to the striped volume. This will be the same for all disks.
6. Click Next.
7. On the Assign Drive Letter or Path screen, choose the desired drive letter for the new volume. Click Next.
8. On the Format Partition screen, choose the option for "Format this volume using the following settings."
9. Choose the NTFS file system. Enter a volume label. Select the option for Perform a Quick Format. Click Next.
10. On the Completing the new Striped Volume Wizard screen, click Finish. The new volume will be formatted and should show a status of Healthy on each of the disks to which it was added.

Windows 7 can create mirrored volumes. Mirroring volumes is done for fault tolerance and redundancy. When a volume is mirrored, a copy of the data written to one volume is also written to a second volume. This way, if one of the volumes becomes corrupted or a disk fails, you can still access your files and data using the copy of the data stored on the other half of the mirror. Follow these steps to create a new mirrored volume:

1. Open the Disk Management MMC snap-in.
2. Right-click on "Unallocated space." Select New Mirrored Volume.
3. The New Mirrored Volume wizard appears. Click Next.
4. The Select Disks screen appears. Select the disks you want to add to the mirrored volume. Click Add.
5. Once the disks are added, you can specify how much disk space from the disks you want to be mirrored. This will be the same for all disks.
6. Click Next.
7. On the Assign Drive Letter or Path screen, choose the desired drive letter for the new volume. Click Next.
8. On the Format Partition screen, choose the option for "Format this volume using the following settings."
9. Choose the NTFS file system. Enter a volume label. Select the option for Perform a Quick Format. Click Next.

10. On the Completing the new Mirrored Volume Wizard screen, click Finish. The new volume will be formatted and should show a status of Healthy on each of the disks to which the mirror was added.

Resize a volume

Sometimes, after you've create a volume, you'll need to change its size. Luckily, once you've created a volume, you're not limited to that size. You can either extend or shrink a volume by taking the following steps:

1. Open the Disk Management MMC snap-in.
2. Right-click on the volume to be extended. Select Extend Volume.
3. The Extend Volume wizard appears. Click Next.
4. The Select Disks screen appears. You can choose to extend the volume on the current disk or extend it to another disk.
5. Click Next.
6. Click Finish. The volume is extended and maintains the same file system as the original volume.

To shrink a volume, do the following:

1. Open the Disk Management MMC snap-in.
2. Right-click on the volume to be shrunk. Select Shrink Volume.
3. When the Shrink Volume window appears, enter the amount of space by which you would like to shrink the volume.
4. Click Shrink. The volume is shrunk and the freed-up space is seen as unallocated space.

Delete a volume

There are several instances in which you may have to remove a volume. For example, you may need to remove a volume if you want to convert a disk to a different format. To delete a volume, do the following:

1. Open the Disk Management MMC snap-in.
2. Right-click on the volume to be deleted and select Delete Volume. You'll receive a warning that all your data will be erased.
3. Click Yes to continue. The volume will be deleted and the freed disk space will be returned to unallocated space.

Next month, I'll cover some more-advanced techniques such as creating a virtual hard disk and advanced diagnostics.



Jorge Orchilles began his networking career as a network administrator for the small private school he attended. He's currently a security operating center analyst, and recently completed his Master of Science degree in management information systems at Florida International University.

©2011 Elsevier Inc. All rights reserved. Printed with permission from Syngress, an imprint of Elsevier. Copyright 2011. "Microsoft Windows 7 Administrator's Reference" (Syngress, 2010) by Jorge Orchilles. For more information on this title and other similar books, please visit elsevierdirect.com.

Related Content

- [Geek of All Trades: Windows 7 Deployment in 7 Easy Steps](#)
- [Windows 7: A Modern Guide to Desktop Deployment](#)
- [Windows 7: The 10 Things to Do First for Windows 7](#)

Active Directory

Protect your Active Directory data

There are a number of tactics for ensuring only the right people have access to the right data within your Active Directory infrastructure.

Darren Mar-Elia

Adapted from "Protecting Critical Data by Managing the Active Directory Identity Lifecycle" (Realtime Publishers)

You must protect your Active Directory-based identity data. It's an important part of ensuring any identity system you put in place that works with Active Directory is protected such that it's able to do its job of authenticating and authorizing the right people to the right resources.

You have to ensure the data within Active Directory is sacrosanct and only users with a business reason to access Active Directory information are granted that access. All the great identity-provisioning processes in the world won't help you if your Active Directory is a free-for-all that anyone can fiddle with to his heart's content. You need to take a deep look at your Active Directory security model and determine the best techniques and best practices for securing the data residing within.

The challenges of securing Active Directory

Managing the Active Directory security model isn't exactly straightforward. The nature of a hierarchical directory service that serves many purposes (including application directory, authentication directory, desktop management directory and so on) means the security model can be a handful. More important, if you don't take a proactive approach to managing your Active Directory security, it can quickly get out of control.

Consider, for example, the simple task of delegating user account management in Active Directory. Because of the granular nature of the Active Directory security model, a seemingly simple task such as managing user accounts could evolve into a dizzying array of permissions you'll have to delegate:

- Permission to create user objects
- Permission to delete user objects
- Permission to move user objects
- Permissions on user object properties (this may break down into sensitive properties, such as department, manager, and group memberships, and non-sensitive properties such as telephone number and office address)
- Permission to reset the user's password or unlock his account
- Permission to control who can change a user's permissions

This list is by no means comprehensive, but it underscores the potential complexity of managing delegation on just this one task. Consider that each of these tasks (or at least groups of them) might be delegated to other subgroups. These permission sets might also vary based on the organizational unit (OU) in which the users are located. Add to the mix that parent objects in Active Directory can inherit permissions from their children (for example, permissions can move from the Marketing OU to the Users OU under Marketing). You can see things can really get gummed up if you're not careful.

Not only is the complexity of the Active Directory security model challenging, it requires discipline to establish a good delegation model and keep it organized over time. One-off requests and unusual business needs drive you to make compromises. The ultimate goal is to protect the data in Active Directory that's critical to your organization's authentication and authorization mechanisms, so it's important to keep a handle on Active Directory security.

Understand the Active Directory security model

Understanding the Active Directory security model is about comprehending how Active Directory is structured. Not unlike a relational database, Active Directory contains a schema that defines the available classes of objects and their associated attributes. A user object in Active Directory is an instantiation of the schema class "user." That user object, as per the schema, contains a set of attributes such as first name, last name, department, manager, phone number and so on.

Each object in Active Directory also has an associated security descriptor. This security descriptor defines the permissions on that object. These show an example of a user object's permission set, or Access Control List (ACL), as viewed from Active Directory Users and Computers.

The ACL is composed of a set of security principals (usually users or groups) that have rights over that object, and the rights or permissions associated with each security principal. A particular permission can be either an "Allow" or "Deny." Allow is the default. This grants a permission to the security principal. If Deny is selected, then that permission is explicitly denied to that security principal. In fact, if an object inherits permissions from its parent and there are clashing Allow and Deny Access Control Entries (ACEs) for a given permission, then typically the Deny will win.

Standard and extended rights

Not every object class in Active Directory has the same set of associated permissions. This is great because it means you can tailor permissions to the type of object involved. Consider this example: a "trigger replication" permission associated with an Active Directory naming context object lets you delegate who can force replication between two domain controllers. Trigger replication has no relevance to a user object, though. In fact, every object has an associated set of "standard rights." These include familiar ones such as:

- Read
- Write
- List
- Create
- Delete
- Read and Write Properties

Beyond the standard rights, a schema class can also have extended rights. Familiar examples of extended rights are the permissions found on a computer object. A computer has permissions such as "Read and Write host name attributes" specific to the computer class of each object.

This extensibility within the Active Directory security model lets you create a rich and granular delegation of tasks for your administrators and users. If you extend the Active Directory schema with a new class of object, it can have its own set of extended rights that control delegation specific to that object type (though you also have to be aware of the various differences across the object classes you want to delegate).

Understand security inheritance

Another aspect that makes the Active Directory security model challenging is the notion of inheriting permissions through the Active Directory hierarchy. A permission set at the top of a domain can trickle all the way down through nested OUs to objects at the lowest levels of the domain hierarchy.

You can control this inheritance, both from the top down as well as the bottom up. For example, say you're setting permissions for user objects within an OU hierarchy composed of a top-level Marketing OU and two sub-OUs called Users-East and Users-West. You want to take advantage of inheritance to set permissions for all user objects at the Marketing OU level and have that trickle down to all user objects in both sub-OUs. You can do so by creating the new ACE within the Marketing OU's ACL (using Active Directory Users and Computers as an example) and then, after setting the permissions, have it apply to all "Descendant User Objects."

If you were running the Users-East OU, you might decide the permissions placed upon your user objects at the higher level don't apply. If you have sufficient permissions, you can essentially turn off the inheritance that comes down from the Marketing OU. You can do so by simply clearing the checkbox in the Advanced section of the ACL editor in Active Directory Users and Computers. That breaks that inheritance chain.

Understand delegation

Within the context of Active Directory, delegation is the process by which you grant permissions to Active Directory objects. This lets users and groups perform specific tasks against Active Directory objects. Delegation implies some kind of orderly plan for giving the right users the right permissions on the right Active Directory objects, across your Active Directory structure.

An example of delegation may be a group called Help Desk Admins, to which all help desk team members belong. You can grant this group the ability to reset user passwords on all user objects within your Active Directory domain. This lets them handle one of the primary tasks of their job. Another common example of delegation is letting admins join computers to the domain. This is a delegated permission on computer objects, typically applied to OUs where computer objects might reside.

Creating a delegation model for Active Directory is probably one of the most important planning tasks you can perform. This is especially true when it involves delegating access to sensitive data held within Active Directory. It doesn't matter whether you've just rolled out Active Directory or if you're in the process of migrating your 10-year-old Active Directory to a new domain structure. In either case, it's never too late to plan and create a delegation model that protects critical objects within Active Directory and delegates access appropriately.

There's a lot of data within Active Directory that can and should be secured, but not all of it relates to your identity system. The following list highlights areas to start with in terms of protecting your identity data within Active Directory:

User properties: Attributes on your users may or may not be sensitive and require delegation. There are certain attributes—such as the user's Job Title, Department and Manager—that are often managed by the HR department. If Active Directory is integrated into an HR system, those attributes might be managed through that system. In that case, you'd want to prevent all users from being able to modify these attributes on their own.

Group memberships: You can use groups to control access to a variety of resources, from public file servers to sensitive database data. Controlling who has permissions to change group memberships is probably one of the first steps you should take in your Active Directory delegation tasks. This translates into who can write to the Members attribute on Active Directory group objects. A user with this permission can modify group memberships.

OU moves: Although moving objects such as users between OUs may seem fairly benign from an identity-management perspective, such a move can often have downstream effects. Some organizations have automated processes associated with OU membership that could change things, such as a user's group memberships or what Group Policy settings they process. These changes could inadvertently grant the user unintended access to resources. As a result, moving OU user objects should be tightly controlled.

Delegation tools

You have some help with respect to delegation. This comes in the form of the Active Directory Users and Computers Delegation of Control (DoC) Wizard. The DoC Wizard is available whenever you right-click a container object (such as a domain or OU) within Active Directory Users and Computers. It essentially turns a set of standard tasks you might want to delegate within Active Directory into a template.

It also lets you create custom delegation tasks by picking object classes and choosing what rights you need on those classes. The DoC Wizard stamps a set of permissions you've requested on the container with which you're working.

The DoC Wizard makes simple delegation tasks easy, but it does have several shortcomings:

- It only supports a small set of delegation tasks (though you can extend the set).
- It's a moment-in-time delegation. In other words, no state of the delegation you just performed is kept. Permissions are simply stamped on the Active Directory objects upon which you're focused. You can't modify that delegation through the wizard after the fact. You'd need to go in and manually edit permissions on the ACL directly.
- The DoC Wizard gives you no bird's-eye view of delegation across your entire Active Directory. Because it's ultimately just stamping privileges, there's no way to "keep track" of what delegations you've done without looking at the ACLs of each Active Directory object of interest.

Active Directory delegation best practices

A number of best practices have emerged over the years that are worth considering as you determine how to create your delegation model and protect your Active Directory data. A common approach is the "role-based" approach to delegating tasks within Active Directory. This involves listing out the tasks you expect people to perform within Active Directory. This list should be fairly long, as you'll want to really flesh out all those things you expect people to do against your Active Directory, especially when it comes to modifying sensitive objects.

The next step is identifying those groups of users that need to perform each task or role. By performing this mapping between groups that need access to Active Directory and the types of access you want to support, you're effectively creating a role-based delegation model you can make real through permissions within Active Directory.

You can design an Active Directory delegation model that breaks security of Active Directory into tasks or roles and then assigns those tasks to user groups. This approach can greatly simplify Active Directory security management and take it from reactive granting of permissions to proactive managing of access to Active Directory data.

Consider all these aspects and how they affect the overall security level of the data stored in your Active Directory directories. Combining and parsing all these techniques can help put your Active Directory data at a safe degree of lockdown.



Darren Mar-Elia is a Microsoft Group Policy MVP, creator of the popular Group Policy site gpoguy.com and coauthor of "Microsoft Windows Group Policy Guide" (Microsoft Press, 2005). He's also CTO and founder of SDM Software Inc. Reach him at Darren@gpoguy.com.

For more on this and other Realtime Publishers titles, check out the [Realtime Publishers Web site](#).

Related Content

- [Identity and Access Management: Access Is a Privilege](#)
- [Identity and Access Management: Filling the Gap in Identity and Access Governance](#)
- [Security: 19 Smart Tips for Securing Active Directory](#)

Security

Don't forget the basics

While much of the focus—and much of the funding—goes to advanced security technologies, it pays to not overlook the fundamentals.

John Vacca

Adapted from “Computer and Information Security Handbook” (Elsevier Science & Technology books).

Many organizations spend a great deal of time and money addressing perimeter defenses. While those are indeed important, organizations can sometimes overlook some fundamental security processes and procedures. Just simple password-management tactics and configuration tricks can go a long way toward increasing your security posture.

Change default account passwords

Nearly all new network devices come preconfigured with a password/username combination. This combination is included with the setup materials and is documented in numerous locations. Very often, these devices are gateways to the Internet or other internal networks.

If you don't change those default passwords upon configuration, it becomes a trivial matter for an attacker to get into these systems. Hackers can find password lists on the Internet, and vendors often include default passwords in their online manuals.

Use robust passwords

With the increased processing power in any PC or laptop and password-cracking software such as the Passware products and the AccessData Password Recovery Toolkit, cracking passwords is fairly simple and straightforward. For this reason, it's extremely important to create robust passwords. Complex passwords are hard for users to remember, though, so the challenge is to create passwords they can remember without writing them down.

One solution is to use the first letter of each word in a phrase, such as “I like to eat imported cheese from Holland.” This becomes lIteicfH, which is an eight-character password using upper and lowercase letters. You can make this even more complex by substituting an exclamation point for the letter l and substituting the number 3 for the letter e, so the password becomes!It3icfH. This is a fairly robust password that a user can easily remember.

Close unnecessary ports

A computer's ports are logical communication access points over a network. Knowing what ports are open on your computers will help you understand the types of access points you have available. The well-known port numbers are 0 through 1023.

Some easily recognized ports and their uses are listed here:

- Port 21: FTP
- Port 23: Telnet
- Port 25: SMTP
- Port 53: DNS
- Port 80: HTTP
- Port 110: POP (Post Office Protocol)
- Port 119: NNTP (Network News Transfer Protocol)

Open ports that aren't necessary can be an entrance into your systems. Open ports that are open unexpectedly could be a sign of malicious software. Therefore, identifying open ports is an important security process. There are several tools that will help you identify open ports. The built-in command-line tool netstat will help you identify open ports and process IDs by using the following switches:

- **a:** Displays all connections and listening ports
- **n:** Displays addresses and port numbers in numerical form
- **o:** Displays the owning process ID associated with each connection

Other helpful tools for port management include CurrPorts, a GUI tool that lets you export the results in delimited format, and TCPView, a tool provided by Microsoft.

Patch, patch, patch

Nearly all OSes have a mechanism for automatically checking for updates. You should make sure this notification system remains turned on. Although there's some debate as to whether updates should be installed automatically, you should at least be notified of updates. You might not want to have them installed automatically, as patches and updates have been known to cause more problems than they solve. However, don't wait too long before installing updates because this can unnecessarily expose your systems to attack. A simple tool that can help keep track of system updates is the Microsoft Baseline Security Analyzer, which also will examine other fundamental security configurations.

Don't use administrator accounts for personal tasks

A common security vulnerability develops when you, as the systems administrator, conduct administrative or personal tasks while logged into your computers with administrator rights. Tasks such as checking e-mail, surfing the Internet, and testing questionable software can expose the computer to malicious software. This also means malicious software may be able to run with administrator privileges, which can create serious problems. To prevent this, you should log into your systems using a standard user account to prevent malicious software from gaining control of your computers.

Restrict physical access

With such a focus on technology, it's often easy to overlook the non-technical aspects of security. If an intruder can gain physical access to a server or other infrastructure asset, the intruder will own the organization. Critical systems should be kept in secure areas. A secure area is one that provides the ability to control access to only those who need access to the systems as part of their job.

A locked room is a good start. Only the server administrator should have a key, and you should store the spare key in a safe somewhere within the executive suites. The room should not have any windows that can open. The room shouldn't even have any labels or signs identifying it as a server room or network operations center. You should most definitely not store the server equipment in a closet where other employees, custodians or contractors can gain access. Review the validity of your security mechanisms during a third-party vulnerability assessment.

Don't forget paper

With the advent of digital technology, some people have forgotten how information was stolen in the past: on paper. Managing paper documents is fairly straightforward. Use locking file cabinets, and make sure they're locked consistently. Extra copies of proprietary documents, document drafts, and expired internal communications are some of the materials that should be shredded. Create a policy to inform employees of what they should and shouldn't do with printed documents.

This example of trade secret theft underscores the importance of protecting paper documents: A company surveillance camera caught Coca-Cola employee Joya Williams at her desk looking through files and "stuffing documents into bags," according to FBI officials. Then in June, an undercover FBI agent met at the Atlanta airport with another culprit, handing him \$30,000 in a yellow Girl Scout Cookie box in exchange for an Armani bag containing confidential Coca-Cola documents and a sample of a product the company had under development.

These fundamental steps toward securing your physical and digital environment are just the beginning. They should provide some insight into where to start building a secure organization.



John Vacca is an information technology consultant, professional writer, editor, reviewer and internationally known best-selling author based in Pomeroy, Ohio. He has authored more than 50 titles in the areas of advanced storage, computer security and aerospace technology. Vacca was also a configuration management specialist, computer specialist, and the computer security official (CSO) for NASA's space station program (*Freedom*) and the International Space Station Program from 1988 until his retirement from NASA in 1995.

For more on this and other Elsevier titles, check out [Elsevier Science & Technology books](#).

Related Content

- [Cloud Computing: Cloud operations and security](#)
- [Utility Spotlight: Microsoft Security Essentials](#)
- [Security Management: The Scary New Hacking Trend](#)

Cloud computing

Challenges in cloud configuration

The selection, configuration and performance of your cloud-based applications will have a massive impact on performance.

Dan Marinescu

Adapted from "Cloud Computing: Theory and Practice" (Elsevier Science & Technology books)

Developing efficient cloud applications comes with many of the same challenges posed by the natural imbalances found in computing, I/O and communication bandwidths of physical systems. These challenges are greatly amplified by the scale of the system, its distributed nature and the fact that virtually all applications are data-intensive.

Though any cloud-computing infrastructure will ideally attempt to automatically distribute and balance processing loads, you're still left with the responsibility of placing the data close to the processing site and identifying its optimal storage strategy. One of the main advantages of cloud computing—the shared infrastructure—can also be a drawback.

Performance isolation is nearly impossible to reach in a real system, especially when the system is heavily loaded. This is even more difficult with cloud computing. The performance of virtual machines (VMs) fluctuates based on the load, infrastructure services, environment and number of users. Security isolation is another challenging factor to identify on multi-tenant systems.

Reliability is also a major concern. You can expect node failures whenever a large number of nodes compete for compute resources. Choosing an optimal instance (in terms of performance isolation, reliability and security) from those offered by the cloud infrastructure is a critical factor. Cost considerations also play a role in the choice of the instance type.

The app experience

Many applications consist of multiple stages. Each stage may involve multiple instances running in parallel on the cloud systems and the communications among them. Thus, efficiency, consistency and communication scalability are major concerns for an application developer. Due to shared networks and unknown topology, cloud infrastructures exhibit internode latency and bandwidth fluctuations that often affect application performance.

Data storage plays a critical role in the performance of any data-intensive application. Organizing the storage, choosing storage location and managing storage bandwidth must all be carefully analyzed for optimal application performance. Clouds support many storage options, including off-instance cloud storage, mountable off-instance block storage and persistent storage for the instance lifetime.

Many data-intensive applications use metadata associated with individual data records. For example, the metadata for an MPEG audio file might include the name of the song, the singer, recording information and so on. Metadata should be stored for easy access and storage should be scalable and reliable.

Another important consideration for application performance is logging. It's a delicate balance. Performance considerations limit the amount of data logging, whereas the ability to identify the source of unexpected results and errors is helped by frequent logging. Logging is typically done using instance storage preserved only for the lifetime of the instance. Thus, you should always take measures to preserve the logs for a postmortem analysis.

Application opportunities

You can divide existing cloud applications into several broad categories: processing pipelines, batch-processing systems and Web applications. Processing pipelines are data-intensive and often compute-intensive applications. These represent a fairly large segment of applications currently running on the cloud. There are several types of data-processing applications:

- **Indexing:** The processing pipeline supports indexing large datasets created by Web crawler engines.
- **Data mining:** The processing pipeline supports searching large collections of records to locate items of interests.
- **Image processing:** A number of companies let you store images on the cloud, such as [Flickr](#) and [Google](#). The image-processing pipelines support image conversion, compression and encryption.
- **Video transcoding:** The processing pipeline transcodes from one video format to another (for example, from AVI to MPEG).
- **Document processing:** The processing pipeline converts large collections of documents from one format to another (such as from Word to PDF) or encrypts the documents. It could also use optical character recognition (OCR) to produce digital images of documents.

Batch-processing systems also cover a broad spectrum of data-intensive applications in enterprise computing. Such applications typically have deadlines. Failure to meet these deadlines could have serious economic consequences. Security is also a critical aspect for many batch-processing applications. A non-exhaustive list of batch-processing applications includes:

- Generating daily, weekly, monthly, and annual activity reports for organizations in retail, manufacturing, and other economic sectors.
- Processing, aggregating, and summarizing daily transactions for financial institutions, insurance companies, and health-care organizations.
- Inventory management for large corporations.
- Billing- and payroll-record processing
- Software development management (such as nightly updates of software repositories).
- Automatic testing and verification of software and hardware systems.

Finally, and of increasing importance, are cloud applications for Web access. Several categories of Web sites have a periodic or a temporary presence, such as the Web sites for conferences or other events. There are also Web sites that are active during a particular season such as the holidays. They might also support a particular type of activity, such as income tax reporting with the April 15 deadline each year. Other limited-time Web sites used for promotional activities “sleep” during the night and auto-scale during the day.

It makes economic sense to store the data in the cloud close to where the application runs. The cost per GB is low and processing is more efficient when the data is stored close to the servers. This could lead to several new classes of cloud-computing applications in the years to come. For example, there could be batch processing for decision support systems and other aspects of business analytics.

Another class of new applications could be parallel batch processing based on programming abstractions. Mobile interactive applications that process large volumes of data from different types of sensors and services that combine more than one data source are obvious candidates for cloud computing.

Science and engineering could greatly benefit from cloud computing because many applications in these areas are compute- and data-intensive. Similarly, a cloud dedicated to education would be extremely useful. Mathematical software such as MATLAB and Mathematica could also run on the cloud.

Application development, selection, configuration and performance tuning all become essential activities when balancing a cloud-computing environment. You’ll have many applications running in many different ways, and that application stack will need some monitoring and maintenance.



Dan C. Marinescu was a professor of computer science at Purdue University from 1984 to 2001. Then he joined the Computer Science Department at the University of Central Florida. He has held visiting faculty positions at the IBM T. J. Watson Research Center, the Institute of Information Sciences in Beijing, the Scalable Systems Division of Intel Corp., Deutsche Telecom AG and INRIA Rocquencourt in France. His research interests cover parallel and distributed systems, cloud computing, scientific computing, quantum computing and quantum information theory.

For more on this and other Elsevier titles, check out [Elsevier Science & Technology books](#).

Related Content

- [Cloud computing: Cloud operations and security](#)
- [Cloud computing: Architecting a Microsoft private cloud](#)
- [Cloud computing: Taxing the cloud](#)

Windows 8

File History explained

File History is a new automated system for continuously protecting your personal files stored in several key locations.

Bohdan Raciborski

Any time your personal files change, there will be a copy stored on a dedicated, external storage device of your choice. File History continuously protects your personal files stored in libraries, desktop, favorites and contacts folders. It periodically (every hour by default) scans the file system for changes and copies changed files to another location. Over time, File History builds a complete history of the changes made to any personal file.

File History was introduced in Windows 8, and gives you a new way to protect your files. It supersedes the existing Windows Backup and Restore features of Windows 7 because it was never a very popular application. This leaves your personal data and digital memories quite vulnerable, as any accident can lead to data loss.

In Windows 8, Microsoft is actively trying to:

- Make data protection so easy any Windows user can turn it on and feel confident his personal files are protected
- Eliminate the complexity of setting up and using backup
- Turn backup into an automatic, silent service that does the hard work of protecting your files in the background without any interaction
- Offer a simple, engaging restore experience that makes finding, previewing and restoring versions of personal files much easier

While designing File History, Microsoft used what it had learned in the past and added requirements to address the changing needs of PC users:

- PC users are more mobile than ever. To address that, File History was optimized to better support laptops that constantly transition through power states, and are being connected and disconnected from networks and devices.
- PC users create more data and are more dependent than ever on data. So not only is the data that's currently on the system drive protected, but also any work and data that was created in the past.

When a specific point in time (PiT) version of a file or even an entire folder is needed, you can quickly find it and restore it. The restore application was designed to offer an engaging experience optimized for browsing, searching, previewing and restoring files.

Setting up

Before you start using File History to back up your files, you'll need to set up a drive to which you will save files. Microsoft recommends you use an external drive or network location to help protect your files against a crash or other PC problem.

File History only saves copies of files in your libraries, contacts, favorites and on your desktop. If you have folders elsewhere you want backed up, you can add them to one of your existing libraries or create a new library.

To set up File History:

- Open the File History control panel applet.
- Connect an external drive, refresh the page, and then tap or click "Turn on."

You can also set up a drive in AutoPlay by connecting the drive to your PC, tapping or clicking the notification that appears, then tapping or clicking "Configure this drive for backup." That's it. From that moment, every hour, File History will check your libraries, desktop, favorites and contacts for any changes. If it finds changed files, it will automatically copy them to the File History drive.

Restoring files

When something bad happens and one or more personal files are lost, the restore application makes it easy to:

- Browse personal libraries, folders and files in a manner similar to Windows Explorer
- Search for specific versions using keywords, file names and date ranges
- Preview versions of a selected file
- Restore a file or a selection of files with a tap or click of a mouse

Microsoft designed the restore application for wide-screen displays and wanted to offer a unique, engaging and convenient way of finding a specific version of a file by looking at its preview.

With other backup applications, you have to select a backup set created on a specific date. Then you have to browse to find a specific folder and then find the one file you need. At this point, however, it's impossible to open the file or preview its contents to determine if it's the correct one. You have to restore the file and if it isn't the right version, you have to start over.

With File History, the search starts right in Windows Explorer. You can browse to a specific location and click or tap on the History button in the Explorer ribbon to see all versions of the selected library, folder or individual file. For example, when you select a Pictures library and click or tap on the History button, you'll see the entire history of the library. When you click on a specific file, you can see the entire history of the selected picture.

You can easily navigate to the desired version by clicking on the Previous/Next buttons or by swiping the screen. Once you've found the version for which you're looking, you can click the Restore button to bring it back. The selected version will be restored to its original location.

Instead of protecting the entire system (the OS, applications, settings and user files), File History focuses on your personal files. That's what is most precious and the most difficult to recreate in case of an accident.

Optimized for performance

In the past, most backup applications used the brute-force method of checking for changes in directories or files by scanning the entire volume. This approach could significantly affect system performance and required an extended period of time to complete. File History, on the other hand, takes advantage of the NTFS change journal.

The NTFS change journal records any changes made to any files stored on an NTFS volume. Instead of scanning the volume, which involves opening and reading directories, File History opens the NTFS change journal and quickly scans it for any changes. Based on this information, it creates a list of files that have changed and need to be copied. The process is quick and efficient.

File History was designed to be easily interrupted and quick to resume. This way, File History can resume its operation without needing to start over when a system goes into sleep mode, when a user logs off, when the system gets too busy and needs more CPU cycles to complete foreground operations, or when the network connection is lost or saturated.

File History was designed to work well on any PC, including small form-factor PCs with limited resources and tablets. It uses system resources in such a way as to minimize the impact on system performance, battery life and overall experience.

File History takes into account:

- If the user is present, meaning logged on and actively using the system
- If the machine is on AC or battery power
- When the last backup cycle was completed
- How many changes have been made since the last cycle
- Activity of foreground processes

Based on all of these factors, which are rechecked every 10 seconds, it determines the optimal way to back up your data. If any of those conditions change, the service makes a decision to reduce or increase quota or suspend or terminate the backup cycle.

Optimized for mobile users

When File History is running, it gracefully handles state transitions. For example, when you close the lid of your laptop, disconnect an external drive, or leave home and take your laptop out of range of the home wireless network, File History takes the appropriate action as follows:

- **Lid closed:** When a PC goes into sleep mode, File History detects the power mode transition and suspends operation.
- **Lid opened:** File History resumes its operation at a priority that ensures files are protected without impacting overall system performance, even for gamers. It also waits for all post-“lid open” activities to complete so the system isn’t affected while it’s coming back out of sleep.
- **Dedicated storage device disconnected:** File History detects that the storage device isn’t present and starts caching versions of changed files on a system drive.
- **Dedicated storage device reconnected:** In the next cycle, File History detects the storage device was reconnected, flushes all versions from the local cache to the external drive and resumes normal operation.

Simplicity and peace of mind

Microsoft designed File History with two objectives in mind: to offer the best possible protection of your personal files and to offer ease of use, simplicity and peace of mind.

If you want to take advantage of File History, you only have to make a few, simple decisions. In most cases, these decisions will be limited to only one: which external drive to use. Windows 8 takes care of the rest. File History operates transparently and doesn’t affect the UX, reliability or performance of Windows in any way. Next month, I’ll cover some of the more advanced backup features of Windows 8 and File History.



Bohdan Raciborski is a principal program manager lead for Microsoft. He has more than 20 years of software development and program management experience. His primary focus is developing software and managing teams in areas ranging from enterprise, high availability, and high-performance storage solutions to consumer products for developed and emerging markets.

Related Content

- [Windows 8: Task Manager retuned](#)
- [Windows 8: Come alive with tiles](#)
- [Windows 8: Identify your unique adoption path](#)

Identity Management

Use two-factor authentication to mitigate fraud

You can—and should—use two-factor authentication with mobile and desktop devices.

Dan Griffin and Tom Jones

At the core of all Web-based transactions is the process of managing fraud. It's essential to balance the inconvenience of user authentication with the risk of being spoofed. To put it another way, when is there sufficient knowledge of the user's identity to allow the transaction to proceed?

Authenticating a user's identity is already a difficult problem in the face of continuously evolving attacks. We're now in the midst of a major paradigm shift. Mobile devices, not desktop computers, will be the platform for most transactions.

Mobile devices create new impediments to securing Internet services. Inconsistent and clumsy keyboards make it inconvenient to enter complex passwords, so most devices offer features to "remember" passwords. Mobile devices are also far more easily lost or stolen than desktop hardware. Plus, mobile applications—including those that may have access to saved passwords—have proven difficult to vet. These factors expose user identities to additional risk.

There's another way to view the proliferation of mobile devices. It's not necessarily a problem, but a solution to previous authentication methods stuck in a losing battle with malware developers and thieves.

The separate communication path—namely, the cellular network—to most mobile devices lets you fortify the barriers faced by any malware. You can use that separate path to bind the user to a specific phone number. Thereafter, the user-to-digits binding (or just having a specific SIM chip) can be reconfirmed as necessary during online transactions. This results in an exponential increase in the number of attempts an attacker would need to mount to get the same level of success.

Phone number as proof

Using phone numbers to reduce online fraud isn't new. Banks have been employing this technique for years. What's new is the ability for any online service, small or large, to easily incorporate this capability.

Microsoft recently acquired PhoneFactor, a mobile phone-based multi-factor authentication solution. It's now available to customers as Windows Azure Active Authentication. With PhoneFactor, when an online service requires two-factor authentication, the phone can provide the second factor in a variety of ways:

- The server makes an automated call back to the phone with a voice message including a one-time secret code. The user then types the code into a Web form to complete the desired online transaction.
- It can also send an SMS message, instead of a voice message, to convey the one-time code and prove user possession of the device.
- A variation of the previous process is to configure the system to let the user respond directly to the SMS (or to a phone tree during a voice call). This prevents having to remember the code while switching from one screen or window to another, and having to do additional typing. Instead, the online Web service is notified asynchronously of the valid SMS response from the user, and the requested transaction, which had been pending, is authorized.
- A PhoneFactor-aware mobile app can receive and respond correctly to an authenticator from the cloud service. Unlike the previous variations, this approach explicitly requires a smartphone app installed on the mobile device.
- You can forward a soft token (such as OAuth) sent to the PC or phone to the service consuming the claim. The PhoneFactor server isn't involved in forwarding the token and doesn't even need to know where it's sent.

Given the assumption that most users already have their phone most of the time, the benefit is multi-factor authentication without the usual incremental hardware (token) cost. This approach can have a dramatic impact on your fraud risk/reward curve.

Now imagine you're a malware writer trying to attack a Web site that requires this type of authentication. Typical attacks focus on the device or the communications path to the device. An attack on the HTTP (TLS) connection is insufficient to get access to the Web service because the authorization code is sent via cell. To be successful, an attack must compromise both channels.

Emerging threats

It's good to understand the threats mitigated by new technology, and how and when new threats are introduced. Fortunately, rootkits have not yet become a factor in phone security. PhoneFactor operates at the application layer and thus offers no mitigation of rootkit threats.

At the app level, adding a second factor with no communications path in common with other factors is significant. It's important to ensure there's no code shared between the paths. That means the user must be involved at some level, either by copying a pin number from one process to another or by accepting a challenge on the phone that is then forwarded using one of the mechanisms described here.

Any interaction that doesn't involve a human in the process is open to electronic attack. On the other hand, new threats to one of the authentication paths don't introduce a greater threat than what's already posed by static passwords.

Because the authorization code is different on every access by every user, a successful app-level attack doesn't permit unfettered access to the Web service by the attacker. Thus, the value of the attack is reduced.

Attackers are interested in the greatest return for the least cost, so using multi-factor authentication encourages them to look elsewhere for easier pickings. By making multi-factor authentication relatively inexpensive to implement, services such as PhoneFactor reduce the chance an attack will succeed.

Fancy toys frequently come with a caveat printed on the outside of the box: "Some assembly required." It's the same with PhoneFactor. To help outline the typical integration requirements, the rest of this article describes solutions using PhoneFactor in two different user scenarios.

Web logon

For interoperable Web logon, the goal is to use PhoneFactor for strong authentication. Then you can represent the user identity with a standard Web token format. You can use PhoneFactor in browser-based Web authentication scenarios by pending the request while the back end contacts the user phone. Once the user authenticates, the Web page refreshes and the user is free to proceed with sensitive transactions.

For this to work, the Web application must trust the PhoneFactor service to authenticate the user. It will have to return some representation of the user identity or other user attributes. Whenever possible, Web applications should be designed to use standards-based token formats, such as Security Assertion Markup Language (SAML) or JSON Web Token (JWT). This helps make them interoperable.

PhoneFactor with Active Directory

PhoneFactor isn't just for Web authentication, though. It can also help make user authentication more secure without too much inconvenience for Windows desktop login. Microsoft has long recognized the need to support alternate authentication methods on Windows.

The Credential Provider API offers this extensibility. To begin the process, the user downloads a Credential Provider (CP) to the device. When the user attempts to log on, the CP sends a message to the back-end authentication service, requesting a PhoneFactor challenge. The CP provides an edit box for the user to type in the code received on his phone.

Once the code is verified, the back-end authentication system takes the additional step of issuing a short-lived public key infrastructure (PKI) certificate for the user. The CP uses the certificate to perform a Kerberos logon. As an additional layer of protection, you can have the certificate private key bound to the Type Parameter Model (TPM) security chip on the client device. It can also be bound to a random PIN generated by the CP. This results in a non-exportable, multi-factor credential that's interoperable with existing applications and hardware.

This approach has the benefit of bridging the gap between PhoneFactor as the new authentication mode and the existing Active Directory infrastructure. There are several components involved in two-factor Windows desktop authentication:

- The user logs on with the existing domain password, in addition to a secret code provided by PhoneFactor. This model uses a trusted Web service to manage the interaction between the client and the PhoneFactor back end.
- The CP acquires a certificate from the Certificate Authority (CA) and logs in to Windows Active Directory. The trusted Web service uses the SAML token from Active Directory Federation Services (AD FS) to acquire the certificate for Kerberos/PKINIT logon.

While this scenario has the user entering the PIN in the domain-joined computer, you can also install a phone app that lets the user accept the request on the phone with a single click.

The benefits of two-factor authentication are well known, but have been difficult to achieve in practice because of the expense in getting the second factor into the hands of users. These are a couple of ways to extend the new Microsoft PhoneFactor technology into the authentication realms of interactive intranet logon, as well as standards-based Web logon.



Tom Jones was the founding chair of the American National Standards subcommittee ASC X9A10 on electronic payments. He has worked within the financial services community with multiple organizations including Electronic Data Interchange (EDI X12) and Accredited Standards Committee X9 Inc. on electronic payments, as well as with First Data Corp., Intel Corp. and Microsoft.



Dan Griffin is the founder of JW Secure Inc. and a Microsoft Enterprise Security MVP. He's the author of the books "Cloud Security and Control" (CreateSpace Independent Publishing Platform, 2012) and "The Four Pillars of Endpoint Security: Safeguarding Your Network in the Age of Cloud Computing and the Bring-Your-Own-Device Trend" (CreateSpace Independent Publishing Platform, 2013). He's also a frequent conference speaker and blogger at jwsecure.com/dan.

Related Content

- [Windows Azure: Authenticate Windows Azure with AD FS](#)
- [Desktop Security: Create Custom Login Experiences with Credential Providers for Windows Vista](#)
- [PhoneFactor](#)

Windows 8

Painless printing

There's an entirely new printer support architecture built in to Windows 8 that will improve support both now and in the future.

Adrian Lannin

Ideally, when you plug a new printer into a Windows machine, it should just work. You shouldn't have to go off and find the right driver. That's one of the big benefits of Windows 8. It abstracts the specific printer from the application, so you don't have to worry about what printer you've installed.

How did my team and I make this happen? We've previously shipped a lot of printer drivers with earlier versions of Windows. Windows Vista had about 4,500 drivers, and Windows 7 had about 2,100. Even though Windows 7 had half as many drivers as Windows Vista, there's a better chance it had a driver for the more popular printers. In Windows Vista, we supported a lot of older, no longer widely used printers, so the relevance of the supported devices supported wasn't as good as in Windows 7.

Windows supports tens of thousands of printer models in total. This includes printers supported by drivers only available via Windows Update or the manufacturers' Web sites. When we see printers that don't work, this is often because the manufacturer has chosen to block the installation. We work with manufacturers to get these packages updated, but it does take time.

When we release a new version of Windows, we take the drivers from the previous version and post them to Windows Update. Even though these devices may drop in popularity, you can still just plug them in and automatically get the device working.

The printer population

People tend to keep printers for five to seven years on average. When we want to add support, we have to think ahead and ask questions like, "What devices are people using? Which were the most popular devices over the past several years? What will be the most popular in the future?"

That last part is tricky because soon after we release new versions or updates to Windows, the printer manufacturers release devices we didn't know about. This means that over time, the set of devices we support in any particular version of Windows becomes stale.

At any given moment, about 100 specific printer models make up about 50 percent of the installed base. If we want to support 75 percent of the models currently in use, we need to support about 300 models.

To get to 95 percent, we need to support more than 1,000 models. The problem is even more difficult because the printers that make up this set of 100, 300 or 1,000 change all the time. The 100 printers that represent 50 percent of the market today aren't the same 100 printers that will represent 50 percent next week or next month, and especially not next year. Every day, many people buy and install new printers.

We've taken a brute-force approach to solving this in the past. We've had representatives from the major printer manufacturers working directly with Microsoft, sitting in offices in Redmond, working to check their source code into Windows. They would create a completely new set of in-box drivers for each new release of Windows. This worked, but wasn't very efficient.

In Windows 8, we took a radically different approach. We stopped shipping lots of printer drivers with Windows. Instead, we built a print class driver framework. This framework is extensible, as it supports printing to existing devices. It also helps manufacturers support new devices, even those that haven't been designed. With the ability to support new and planned printers, the number of printers the Windows 8 print class driver framework supports will actually increase over time.

With a print class driver framework, we can get closer to providing a driverless printing experience. You won't have to actually go and find a driver. Instead, the printer just works with the Windows printing system. A true driverless printing experience requires changes to how most printers are designed. The print class driver framework supports this idea, but it's also important to provide as much support for existing devices as possible.

Resource reduction

Besides making great progress in increasing the number of supported devices, we've also been able to reduce the amount of resources required. First, we reduced the amount of disk space needed to support printers and imaging devices from 768MB in Windows Vista to about 184MB in Windows 8. This is an average across different editions and architectures of Windows 8. We've also increased in the relevance of the devices supported directly by Windows.

This is a huge improvement in Windows 8. This reduction in space translates directly into more available storage space for users of hardware with limited storage capacity, which we expect will be a characteristic of some Windows RT computers.

The Windows 8 printer driver model helps us focus our manufacturing partners on a set of code that won't change as much from one version of Windows to the next. We'll be able to more usefully spend those resources on improving quality and performance, instead of constantly repopulating the driver set.

Print class driver architecture

Besides creating an architecture that supports the needs of Windows Store apps, we wanted to ensure the model would also work with existing devices. It had to use technologies familiar to printer manufacturers, so it would be easy for them to implement the new driver technology.

A printer driver does several key things to begin the printing process:

- Configuration lets you change settings, translating the intent to, for example, turn on double-sided printing into the specific command the printer needs to be able to do this. You can adjust configuration options through the UI.
- Rendering translates printed content from the format the Windows print system uses into a format the printer understands. In some cases, the printer may directly understand the native Windows print format (XPS). For those devices, there's no work to do here, unless a users wants to do extra rendering (doing multiple pages per physical sheet of paper is an example of this case). The part of the driver that does rendering is called the render filter.
- The printer informs the user something has happened with event notifications—a job is complete, there's been a paper jam or the printer is out of ink.

Configuration interface

One big change between the old driver model and the Windows 8 driver model is how the interface is provided. In the old printer driver model, the configuration UI was built in to the driver. In the Windows 8 driver model, the manufacturer's UI is completely separate from its driver. Windows 8 will automatically show you the correct type of UI.

This is a better architecture for many reasons: The UI to control the printer is now an app you can invoke when printing from Windows Store apps or Windows desktop apps. This helps printer manufacturers present you with a much richer experience.

If the manufacturer hasn't provided a configuration UI for its device, then Windows provides a standard UI you can use with any printer. However, when the printer manufacturer has decided to invest in providing a customized experience for its device, it can provide an app that replaces the standard Windows UI. Then, when you decide to alter the device configuration or when the device configuration changes during printing (such as during a paper jam), Windows will display the manufacturer's customized app instead.

Rendering

One of the most important functions of a printer driver is to take the content the app produces when you ask it to print and convert it into something the printer can understand. This was one of the most challenging areas of building the Windows 8 print class driver.

Apps like Word or Photoshop use graphics commands to draw content onto the screen or the printer. When they do this, the print system receives the content and converts it to XPS if necessary. It then calls the printer's driver (or, more specifically, the render filter part of the driver) to convert the content to the correct format. This is sent to the printer and your file is printed out.

Probably one of the largest challenges in supporting a wide range of printers is dealing with rendering. Some of the more expensive printers support standard page description languages (PDLs) such as PostScript, Printer Command Language (PCL) and XPS. Less-expensive, consumer-focused devices are manufactured with cost savings in mind, and many of these support proprietary methods of sending the page information to the printer.

Some manufacturers only have a few languages they use across their product line. Others may tweak the language from one model to the next, trying to get the most out of their printer hardware. This leads to a 1-1 mapping between printer driver and printer hardware.

Imagine each PDL as a complete printer driver. It's easy to see that increasing support involves a steadily increasing number of drivers. This is a bit of a simplification. It's possible to create a driver that supports a number of devices. We've often seen drivers that support a series of printers. The key point here is that Windows 7 and earlier versions of Windows didn't do anything to support this design approach.

The printer driver model in Windows 8 supports the idea that a PDL (or driver) can be associated with multiple devices. We've been working with our printer manufacturer partners to have them include an identifier in their devices that generically describes how they're supported. We call this a compatible ID. If a device has a compatible ID that says the device supports XPS, then the print system knows it doesn't need to find a model-specific driver for that device. It can just install a generic XPS driver for the device.

Windows understands the device is a generic XPS printer and can treat it that way. Of course, Windows also understands it's a Fabrikam 1000 printer (or whatever), so if there's a model-specific driver, then Windows will install it. If there's no driver available, Windows can still print to the printer using the class driver.

So, in this example, there's a set of render filters as part of the class driver model. You can install these for any device that implements a matching compatible ID. The logical extension of this idea is that it's quite possible for future devices to be compatible with the print class driver in Windows 8.

We've been working with the printer manufacturers, and they all plan to implement compatible IDs in their devices. Because of this, the number of printers supported in Windows 8 will increase over time. More people will get the experience of being able to use their printers instantaneously from Windows 8 without the need to go and find a driver.

What about all the devices that have proprietary rendering languages? The print class driver supports that model too, but with the disadvantage of needing a separate rendering filter for each small set of models that speak each unique language. There's no way around this. In Windows 8, we've taken a number of filters that address a set of popular models. However, once again, we've been working with the printer manufacturers to improve this position. We expect to see manufacturers produce printers that can more easily use the class driver in the future.

Printing from Windows RT

The reduction in the resources used by the print class driver contributes directly to a smaller footprint for Windows. This is especially valuable on Windows RT. The version 3 printer driver architecture was highly extensible and had evolved over many years into a model that encourages the development of large, complex printer drivers. Some drivers install services that run all the time. This can exhaust battery power and waste processor time.

The need to support printing in Windows RT, and a general desire to make printing more efficient, led us to develop an architecture that more tightly controls what the driver can do. The UI portion of the print experience is now a completely separate component—an app instead of part of the driver. This means that it's also optional. Drivers will work well with the standard Microsoft printing UI. We've also simplified the driver architecture to be more power-efficient by removing service dependencies and reducing the likelihood that additional software will be included with the driver.

With the Windows 8 driver model, we made significant changes to how printer drivers are installed. In Windows 7 and earlier versions of Windows, all printer drivers are stored in the Driver Store. When you plugged in a printer, we'd find the correct driver in the driver store and copy it to a special location where the spooler could use it with your printer. In Windows 8, we've eliminated this extra copying, which removed quite a bit of disk I/O. The print spooler now just knows how to find the driver in the driver store.

The Windows 8 printer driver architecture is a big step forward. It provides good support for a lot of the printers that people already own. It will also support future devices with a small, fast, built-in class driver framework. The performance is great and the disk footprint is small.



Adrian Lannin is the lead program manager for Printing and Scanning in Windows, and for Windows To Go. Within these areas, he's responsible for designing and shipping features for enterprise users as well as home users. Lannin has more than 25 years of experience in the imaging industry.

Related Content

- [Windows 8: Go mobile with Windows 8](#)
- [Windows 8: Identify your unique adoption path](#)
- [Windows 8: The print system reimagined](#)

SharePoint

Secure SharePoint content

Keeping the content stored within your SharePoint storage resources secure might be more complicated than you thought.

Dan Sullivan

Adapted from "The Essentials Series: Securing SharePoint Content" (Realtime Publishers)

Microsoft provides SharePoint-specific client applications such as SharePoint Workspace, but the standard for SharePoint access is a browser. There are many advantages to using a browser with SharePoint, including access to SharePoint hosting services and cross-platform support.

With these advantages, however, come technical challenges. Together, these challenges create conditions that hackers can exploit to steal or leak confidential data stored in SharePoint repositories.

You almost certainly have valuable content in your SharePoint collaboration sites. The documents stored there might include trade secrets about business processes and products, strategic plans for business expansion, or confidential information about clients and customers. The specifics aren't important. The crucial point is that your SharePoint content will be targeted by attackers.

Think like an attacker

When you try to understand the risks to your content, it helps to think like an attacker. You don't have legitimate access to the corporate network. You can't easily walk the halls of the company. Your best bet is to steal information electronically.

You could invest the time and effort to probe the corporate network for vulnerabilities that would help you gain access to servers and applications. You could then work to avoid detection as you probe servers for software vulnerabilities, run password-detection programs, and try other means of compromising the servers.

But that's just one approach to stealing content. There's another path with potentially less resistance: targeting client devices. Client devices are often subject to fewer controls than servers. Users install applications or browser extensions. Laptops and mobile devices aren't always connected to secure corporate networks. Users might browse compromised sites or open malicious e-mails that lead to malware infection.

In spite of your best efforts, client devices such as desktops and laptops can harbor malware that makes them easier targets than hardened servers. There are two particular types of external threats relevant to protecting SharePoint content: malware on the endpoint and browser cache mining.

Endpoint malware

There are many forms of malicious software (malware), including Trojan Horses, keyloggers and rootkits. Trojan Horses are applications that appear benign, but are in fact malicious. For example, a utility for displaying weather updates on your desktop might also scan your drives and copy data to a centralized server.

Keyloggers are designed to capture keystrokes as you type. This helps attackers collect usernames and passwords. Of course, attackers will also end up collecting everything else you type, but text-processing tools can easily analyze large volumes of keystroke data to find information of particular interest to attackers.

Rootkits are sets of programs that attack the OS at low levels. This helps them circumvent OS security controls. This kind of malware is especially difficult to eliminate. These are just three types of malware you might find on client devices.

Browser cache mining

A particularly promising target for some attackers is your browser cache. The browser cache is used to temporarily store data to improve browser performance. A cache is quite useful when navigating Web pages. For example, assume you're reading a long article divided into several Web pages. If at some point during the article you decide to go back to the previous page, you'd probably click the previous page button on your browser. In theory, your browser could simply download the page again, but doing so would take time and consume network resources. Instead, your browser keeps copies of data in local temporary storage known as the cache.

This setup sounds like a reasonable resource trade-off. A certain amount of local storage is dedicated to temporarily storing browsing data in order to save time and network resources. Here again, you have to think like an attacker. Your SharePoint content, which may be well protected on the SharePoint server, might also be cached on endpoint devices. Malware designed to scan and analyze browser caches could gain access to your SharePoint content. This type of attack is known as cache mining.

The benefits of caching apply equally well to SharePoint and other Web content, so you won't necessarily want to eliminate caching. You do, however, want to ensure confidential or sensitive information isn't retained in the cache longer than needed. It should be cleared after a user's SharePoint session has ended.

In addition to caching data to improve browser performance, some online applications let you store copies of data on client devices so those applications can work offline. This is useful for users who work in multiple locations or find themselves traveling with intermittent Internet access. The more data that's stored locally, however, the greater the potential for a data leak.

Internal threats: careless and malicious employees

There are more pedestrian methods an attacker might use to steal valuable information. Some threats to business information stem from carelessness, while others have more malicious origins.

During the course of the day, many of us try to streamline tasks to save time. You might download documents and e-mail copies to collaborators who don't have access to the content in SharePoint. You might justify this by thinking you're working more efficiently. Your business partner might need a piece of information in a document, and you get it to him in the most efficient way possible. The problem with this approach is that you lose control of that digital copy of the document once it's e-mailed.

At that point, you have no control over how your collaborator shares the document. Will it be deleted after only the necessary information is reviewed? Will it be forwarded to someone else? Will copies be stored on e-mail servers that might be attacked in search of valuable data? There's also the question of whether additional content in the document needs to be shared. Does the collaborator need to know everything in the document? Could you compromise your business by sharing too much information?

These are difficult questions to answer. Rather than risk the negative consequences of such carelessness, organizations can implement security controls that block inappropriate copying of SharePoint content from client devices. These controls have the additional benefit of blocking disgruntled employees or others who might intentionally attempt to steal data such as client lists or design documents.

The bring your own device movement

Organizations are grappling with the increase of bring your own device (BYOD) practices. Employees are working with their own laptops, tablets and smartphones to access content and applications on corporate networks. Although there are tools to help manage laptops and mobile devices, organizations have less control over endpoints.

There are limits to what an IT organization can impose on BYOD users. This reality creates potential conflicts. IT professionals may be responsible for protecting corporate information assets, but employees expect reasonable control over their devices. The potential for conflicting expectations is high.

Consider a user who installs a browser extension to help monitor prices. When the user searches for a product or service, such as a flight from San Francisco to New York, the browser extension checks multiple sites and displays information about the best prices. To perform this kind of service, the extension needs access to the browser data. An employee might be willing to allow access to personal browsing information, but what about work-related browsing? Do IT professionals responsible for information security want SharePoint content exposed in that way?

In addition to apps and extensions your users intentionally install, you need to consider the possibility that employee-owned devices might not have adequate security controls in place. For example, do employees:

- Keep their anti-malware software up-to-date? Databases used by anti-malware software are updated frequently. Endpoint anti-malware programs should be configured to automatically check for and download data and program updates.
- Run vulnerability scanners to check for known vulnerabilities in software installed on the client? Vulnerability scanners were once limited to network and server administrators, but even end users can run tools such as Microsoft Baseline Security Analyzer on desktop and laptop devices.
- Update OS and application software? Malicious software can take advantage of vulnerabilities in widely used productivity or utility software. Keeping software up-to-date is a key security practice.

BYOD has many advantages for both employees and employers. As with so many technologies and practices, though, there are benefits and drawbacks to employee-owned devices being used for business purposes.

Securing SharePoint content is challenging enough. Sending copies of documents from well-secured SharePoint servers to poorly secured endpoint devices can leave your information vulnerable to theft or leakage. Threats range from malware to malicious employees. Changes in the way you work, particularly the increasing practice of BYOD, compound the SharePoint security challenges you already face.



Dan Sullivan has more than 20 years of IT experience in application design, systems architecture and enterprise security. He's written and presented extensively about systems architecture, infrastructure management, and aligning business and IT strategies. He's written several books, including "The Shortcut Guide to Prioritizing Security Spending," "The Definitive Guide to Security Management" and "The Definitive Guide to Information Theft Prevention," all from Realtime Publishers.

For more on this and other titles from Realtime Publishers, check out the company's [Web site](#).

Related Content

- [SharePoint 2010: Grow your SharePoint farm](#)
- [SharePoint 2013: Add power to Office apps](#)
- [Microsoft SharePoint 2010: Mandatory Services in SharePoint](#)

Windows Confidential

Respect the brand

The Windows 95 10K Program goes on tour to promote Windows 95 (and silences an annoying IT guy).

Raymond Chen

During the development process for Windows 95, we set up what we called the 10K Program. This program sent select Windows 95 team members to a variety of mostly large companies.

The purpose was to introduce the OS that was then under development to the IT departments at those companies. They would show off the new features and demonstrate how it would boost employee productivity. Perhaps a more important purpose was for the Windows team members to learn about ways they could improve Windows 95 to make it easier for IT departments to deploy and manage within their organizations.

Instead of holding meetings or giving presentations, these 10K visits followed a highly practical format. They were set up like a pilot program. The IT departments would prepare for the visit by selecting a number of employees to participate. The 10K team would then walk through installing Windows 95 on their machines. Then they would boot up and see what happened.

Did they still have access to the corporate network? Did the company's proprietary line-of-business (LOB) software still work? Could they still print their budget reports? Did Windows 95 make it easier to find and configure a color printer?

The 10K team would stay on-site for a few days to assist with initial deployment and observe the test subjects as they settled in with the new OS. Then they'd leave a set of installation CDs behind so the IT department could expand its testing after the team left.

Coke, no Pepsi

It was considered quite a privilege to be selected to go on one of these trips. Not only were you able to get out of the office for some travel, but you also got to spend time talking to actual customers about the project you'd been working on so hard for so long. The 10K Program members would follow up periodically after their site visits in order to get feedback on the OS after users spent more time with it and after the IT departments had an opportunity to test it more widely within their organization.

As you might imagine, these on-site visits provided extremely valuable feedback to the Windows 95 team. We learned about all sorts of unusual hardware, software, network and printer configurations, as well as learning about how people used Windows in their everyday work. The team members also got a chance to offer simple tips and tricks to the employees as they observed their work processes, often basic shortcuts such as, "Click the Printer icon on the toolbar if you want to print the document."

The name 10K came from the ostensible target of upgrading 10,000 computers to Windows 95 through this program. I don't know whether they actually hit the goal, but the number 10,000 was really just part of the name. The purpose of the program was not to hit the 10,000 target, but rather to introduce Windows 95 to the IT departments and return with valuable information. The number 10,000 was largely arbitrary.

One of these 10K visits was to the Coca-Cola Company. The 10K team was greeted by representatives from the IT department. They were chatting as they made their way through the lobby, down the hall, up the elevator, and then over to the lab rooms that had been set up. One of the IT people brandished an OS/2 CD and wouldn't stop challenging the 10K team. "OS/2 this, OS/2 that, what do you think of this, how are you going to respond to that?"

Eventually, the 10K team leader politely made the following offer: "I'll make you a deal. If you don't talk about OS/2, then I won't talk about Pepsi." There was no mention of OS/2—or Pepsi—for the remainder of the visit.



Raymond Chen's *Web site, The Old New Thing, and identically titled book (Addison-Wesley, 2007) deal with Windows history, Win32 programming, and what computer programmers and fashion models have in common.*

Related Content

- [Windows Confidential: Retire the rubber stamp](#)
- [Windows Confidential: The feature battle](#)
- [Windows Confidential: Signs your project is doomed](#)

Learn about the [Use of Microsoft Copyrighted Content](#)