

Microsoft®

MCTS EXAM

70-680

Configuring Windows® 7

**UPDATES
INSIDE**



Ian McLean and
Orin Thomas

SELF-PACED

Training Kit

Book Update/Page Corrections

For ISBN: 978-0-7356-2708-6

To learn more about the complete book, visit <http://oreilly.com/catalog/9780735627086/>

PUBLISHED BY

Microsoft Press

A Division of Microsoft Corporation

One Microsoft Way

Redmond, Washington 98052-6399

Copyright © 2011 by Orin Thomas and Ian McLean

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Information in this document, including URL and other Internet Web site references, is subject to change without notice.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without express written permission of Microsoft Corporation.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Introduction

This training kit is designed for IT professionals who operate in enterprise environments that use Windows 7 as a desktop operating system. You should have at least one year of experience in the IT field, as well as experience implementing and administering any Windows client operating system in a networked environment.

You should be able to install, deploy, and upgrade to Windows 7, including ensuring hardware and software compatibility. Additionally, you should be able to configure preinstallation and postinstallation system settings, Windows security features, network connectivity applications included with Windows 7, and mobile computing. You should also be able to maintain systems, including monitoring for and resolving performance and reliability issues and have a basic understanding of Windows PowerShell syntax.

By using this training kit, you will learn how to do the following:

- Install, upgrade, and migrate to Windows 7.
- Deploy Windows 7.
- Configure hardware and applications.
- Configure network connectivity.
- Configure access to resources.
- Configure mobile computing.
- Monitor and maintain systems that run Windows 7.
- Configure backup and recovery options.

MORE INFO FIND ADDITIONAL CONTENT ONLINE

As new or updated material that complements your book becomes available, it will be posted on the Microsoft Press Online Windows Server and Client Web site. The type of material you might find includes updates to book content, articles, links to companion content, errata, sample chapters, and more. This web site will be available at <http://www.microsoft.com/learning/en/us/training/format-books-support.aspx#tab2> and will be updated periodically.

Lab Setup Instructions

The exercises in this training kit require a minimum of two client computers or virtual machines running Windows 7 Enterprise or Ultimate editions. Instructions for configuring the first of these computers are given in Chapter 1, “Install, Migrate, or Upgrade to Windows 7.” Instructions for configuring the second of these computers are given in Chapter 6, “Network Settings.” You need an additional hard disk (internal or external), formatted using the NTFS file system, installed on the first of these computers.

You can obtain an evaluation version of Windows 7 Enterprise from the Microsoft Download Center at the following address: <http://technet.microsoft.com/en-us/evalcenter/default.aspx>.

All computers must be physically connected to the same network. We recommend that you use an isolated network that is not part of your production network to do the practices in this book. To minimize the time and expense of configuring physical computers, we recommend you use virtual machines. To run computers as virtual machines within Windows, you can use Hyper-V, or third-party virtual machine software.

Hardware Requirements

You can complete almost all the practice exercises in this book using virtual machines rather than real hardware. The minimum and recommended hardware requirements for Windows 7 are listed in Table 1.

TABLE 1 Windows 7 Minimum Hardware Requirements

HARDWARE COMPONENT	MINIMUM REQUIREMENTS	RECOMMENDED
Processor	1 GHz 32-bit (x86) or 64-bit (x64) processor	2 GHz or faster
RAM	1 GB	2 GB
Disk space	40 GB	60 GB
Graphics adapter	Supports DirectX 9 graphics Has a Windows Display Driver Model (WDDM) driver Pixel Shader 2.0 hardware 32 bits per pixel 128 MB graphics memory	As minimum requirement, but with 256 MB of graphics memory

Windows 7 Starter

Windows 7 Starter is available from retailers and on new computers installed by manufacturers. It does not support or include the Windows Aero user interface, DVD playback, Windows Media Center, IIS Web Server, or Internet connection sharing. You cannot join a computer with this edition of Windows to a domain. This edition does not support enterprise features such as Encrypting File System (EFS), AppLocker, DirectAccess, BitLocker, Remote Desktop Host, and BranchCache. This edition supports a maximum of one physical processor.

Windows 7 Home Basic

Windows 7 Home Basic is available only in emerging markets. It does not support or include the Windows Aero user interface, DVD playback, Windows Media Center, or IIS Web Server. You cannot join a computer with this edition of Windows 7 to a domain. This edition does not support enterprise features such as EFS, AppLocker, DirectAccess, BitLocker, Remote Desktop Host, and BranchCache. This edition supports a maximum of one physical processor. The x86 version supports a maximum of 4 GB of RAM, whereas the x64 version supports a maximum of 8 GB of RAM.

NOTE MULTIPROCESSOR AND MULTICORE

Windows 7 Starter, Home Basic, and Home Premium will recognize only one physical processor. Other editions support two physical processors. All 32-bit versions of Windows 7 can support up to 32 processor cores, while 64-bit versions can support up to 256 processor cores.

Windows 7 Home Premium

Windows 7 Home Premium is available from retailers and on new computers installed by manufacturers. Unlike the Starter and Home Basic editions, the Home Premium edition supports the Windows Aero UI, DVD playback, Windows Media Center, Internet connection sharing, and the IIS Web Server. You cannot join this edition of Windows 7 to a domain, and it does not support enterprise features such as EFS, AppLocker, DirectAccess, BitLocker, Remote Desktop Host, and BranchCache. The x86 version of Windows 7 Home Premium supports a maximum of 4 GB of RAM, whereas the x64 version supports a maximum of 16 GB of RAM. Windows 7 Home Premium supports one physical processor.

Windows 7 Professional

Windows 7 Professional is available from retailers and on new computers installed by manufacturers. It supports all the features available in Windows Home Premium, but you can join computers with this operating system installed to a domain. It supports EFS and Remote Desktop Host but does not support enterprise features such as AppLocker, DirectAccess, BitLocker, and BranchCache. Windows 7 Professional supports up to two physical processors.

computer. This type of installation is suitable when you are deploying Windows 7 to a small number of computers.

- **Unattended installation** You can perform an unattended installation of Windows 7 by using an installation file called Unattend.xml. These installation files store answers to the questions asked by the Setup Wizard. When the Windows 7 installation process starts, Windows checks for attached USB storage devices that have this file in their root directory. Unattended installations are suitable when you need to deploy Windows 7 to a large number of computers because you do not have to interact with them manually, responding to prompts, as the installation progresses.

Clean Installations

A clean installation is one performed on a computer that does not currently have an operating system installed. This might be a brand-new computer arriving straight from the factory, or it might be an older computer with a brand-new hard disk drive on which you wish to install Windows 7. You can use any installation media to perform a clean installation. Although you perform a clean installation in the practice exercise at the end of this chapter, the next few pages will describe in detail some of the options and concepts that you encounter during the installation process.

The first page that you encounter when performing an installation asks you which language you wish to install, the time and currency format, and what keyboard or input method you are using for the installation. This selection is important, not only because installing Microsoft Windows in a foreign language can be a challenge if you do not understand that language, but because even keyboards from other English-speaking countries have layouts that differ from the standard layout of a U.S. keyboard. If you are installing Windows 7 for users that need access to multiple keyboard layouts, you can add these alternate layouts once the installation process has completed. Then the user can switch between them as necessary.

The next page, shown in Figure 1-1, is the Install Windows 7 page. From here, click Next and then you can start the installation by clicking Install Now. You can also access some of the Windows 7 repair tools by clicking Repair Your Computer. You will learn more about the repair options in Chapter 14, "Recovery and Backup." Clicking the What To Know Before Installing Windows option provides you with general installation advice, such as ensuring that your computer needs to meet the minimum hardware requirements and that you should have your product key ready.

The next step is to review and accept the Windows 7 license terms. This is followed by choosing what type of installation you want to perform: Upgrade or Custom (Advanced). When you are performing a clean installation, you should select Custom (Advanced). Almost all installations of Windows 7 that you will perform will be of the Custom (Advanced) type rather than upgrades. You can initiate upgrade installations only from within Windows Vista or Windows 7. You learn more about upgrading to Windows 7 in Lesson 2.



FIGURE 1-1 The Install Windows 7 page

The next step in the installation process is determining where the Windows 7 files will be stored. Windows 7 needs a minimum of 15 GB of free space, though you should generally allocate more than this amount. From this page, it is possible to partition an existing disk into smaller volumes. You can do this by clicking Drive Options (Advanced). The installation routine recognizes most Integrated Drive Electronics (IDE), Serial Advanced Technology Attachment (SATA), and Small Computer System Interface (SCSI) disk drives automatically. If your computer has special disk drive hardware, such as a redundant array of independent disks (RAID) array, it may be necessary to use the Load Driver option. It is necessary to use this option only if the disk that you want to install Windows on is not shown as a possible install location. If your disk is shown as an available option, Windows 7 has already loaded the appropriate drivers. Once you select the location where you want to install Windows 7, the Windows 7 installation process begins.

NOTE INSTALLING TO VHD

When performing a clean install of Windows 7 Enterprise and Ultimate editions, you have the additional option of installing to a virtual hard disk (VHD) file rather than directly to the volume. You will learn more about the steps required to do this in Chapter 2.

After the computer reboots, you need to specify a user name and a computer name. The user name you specify is the default administrator account for that computer. The account named Administrator, used in previous versions as the default administrative account, is disabled by default. It is possible to enable this account only by modifying Group Policy. Because the user name that you specify during setup is the default administrator account for that computer, organizations that are performing Windows 7 deployment need to come

dual-boot with, you may end up wiping that volume and replacing it with a fresh installation of Windows 7. When the installation completes, you can dual-boot between operating systems and Windows 7 is the default operating system. You will learn about configuring the default operating system later in this lesson.

Dual-Boot and Virtual Hard Disks

An exception to the rule that you require a separate partition for each operating system if you want to dual-boot is booting from a VHD file. You can install and then boot Windows 7 Enterprise and Ultimate editions, as well as Windows Server 2008 R2, from VHD files. It is possible to boot from VHD only on computers that have a Windows 7 or Windows Server 2008 R2 boot environment. This means that you cannot dual-boot Windows XP or Windows Vista computers with Windows 7 installed on a VHD file, though it is possible to triple-boot with Windows 7 to VHD if you already have a computer that dual-boots between an earlier version of Windows and Windows 7. This is because a computer that you have configured to dual-boot to Windows 7 already has the Windows 7 boot environment present. You will learn how to install Windows 7 onto a VHD file in Chapter 2.

Configuring the Default Operating System

When you configure a computer to dual-boot, one of the operating systems is selected as the default. This means that the computer boots into this operating system by default unless the user intervenes to select the other operating system. To change the default operating system using the graphical user interface (GUI), perform the following steps:

1. From the Start menu, open the Control Panel. Select Small Icons from the View By drop-down list.
2. Click System and then click Advanced System Settings. This opens the System Properties dialog box.
3. On the Advanced tab, click Settings in the Startup And Recovery section. This opens the Startup And Recovery dialog box, shown in Figure 1-7.
4. From the Default Operating System drop-down list, select which operating system is booted by default.

To configure the default operating system using the Bcdedit.exe command-line utility, perform the following steps:

1. Open an administrative command prompt by right-clicking Command Prompt in the Accessories folder of the Start menu and choosing Run As Administrator. Click OK when prompted by the User Account Control dialog box.
2. Enter the command **bcdedit /enum** to view a list of the current boot menu entries, similar to what is displayed in Figure 1-8.
3. To change the default entry, use the command **bcdedit /default** and then list the identifier. In Figure 1-8, this would be {878f7be0-2163-11de-b92d-d86aaca536b1}.

12. On the Select Your Computer's Current Location page, click Home Network.
13. If you are presented with the Do You Want To Create A Homegroup? page, shown in Figure 1-13, select all libraries and then make a note of your Homegroup password. Click Next. After you perform this step, the installation is complete. Turn off the computer.




FIGURE 1-13 Create Homegroup

Lesson Summary

- Windows 7 comes in six different editions. Each edition has a different set of features. Only the Professional, Enterprise, and Ultimate editions can be configured to join a domain. Only the Enterprise and Ultimate editions include enterprise features such as BitLocker, AppLocker, DirectAccess, and BranchCache.
- Windows 7 comes in two different versions. The 32-bit version is also known as the x86 version. The 64-bit version is also known as the x64 version. The 32-bit versions support a maximum of 4 GB of RAM. The x64 versions support between 8 GB and 128 GB, depending on the edition.
- Windows 7 Home Basic and Starter editions require a 1 GHz x86 or x64 CPU, 512 MB of system memory, 20 GB HDD, and a 32-MB graphics adapter that supports DirectX9.
- Windows 7 Home Premium, Professional, Enterprise, and Ultimate editions require a 1 GHz x86 or x64 CPU, 1 GB of system memory, 20 GB HDD, and a 128-MB graphics adapter that supports a WDDM driver, Pixel Shader 2.0, 32 bits per pixel, and DirectX9 graphics.




FIGURE 1-14 Windows Anytime Upgrade

Prior to attempting to perform the upgrade from Windows Vista to Windows 7, you should run the Windows 7 Upgrade Advisor. The Windows 7 Upgrade Advisor is an application that you can download from Microsoft's Web site that will inform you if Windows 7 supports a computer running the current hardware and software configuration of Windows Vista. Prior to running the Windows 7 Upgrade Advisor, you should ensure that all hardware that you want to use with Windows 7, such as printers, scanners, and cameras, are connected to the computer. The Upgrade Advisor generates a report that informs you of which applications and devices are known to have problems with Windows 7. A similar compatibility report is generated during the upgrade process, but the version created by the Windows 7 Upgrade Advisor is more likely to be up to date.

MORE INFO WINDOWS 7 UPGRADE ADVISOR

You can obtain the Windows 7 Upgrade Advisor from the following Web site:
<http://www.microsoft.com/windows/windows-7/upgrade-advisor.aspx>.

NOTE USE A SEARCH ENGINE

Another way to determine whether a particular hardware device or application is compatible with Windows 7 is to use a search engine. It is likely that someone before you has attempted to use the set of hardware devices and applications that you are interested in with Windows 7. If they have had a problem with it, it is likely that they have posted information about it in a support forum, a blog, or somewhere else on the World Wide Web.

Windows Vista and Windows 7 have the same basic hardware requirements. This means that you should not have to upgrade the RAM or processor for your computer running Windows Vista to support Windows 7, though you need at least 10 extra gigabytes on the Windows Vista volume to perform the upgrade.

You should keep the following in mind prior to and during the upgrade from Windows Vista to Windows 7:

- Perform a full backup of the computer running Windows Vista prior to performing the installation. That way, if things go wrong, you can do a full restore back to Windows Vista.
- You must ensure that Windows Vista has Service Pack 1 or later installed before you can upgrade it to Windows 7.
- Ensure that you have the Windows 7 product key prior to the upgrade.
- You cannot upgrade between processor architectures. An x86 version of Windows Vista cannot be upgraded to an x64 version of Windows 7, and vice versa.
- You can upgrade only to an equivalent or higher edition of Windows 7. You can upgrade Windows Vista Home Premium to Windows 7 Home Premium, Professional, Enterprise, or Ultimate, but not to Windows 7 Starter. Windows 7 Professional is equivalent to Windows Vista Business. It is only possible to upgrade to Windows 7 Enterprise edition from Windows Vista Business or Enterprise edition. Windows Vista Enterprise edition can only be upgraded to Windows 7 Enterprise edition.
- Ensure that there is at least 10 GB of free disk space on the Windows Vista volume prior to attempting the upgrade.

You will perform an upgrade from Windows Vista to Windows 7 in the practice exercise at the end of this lesson.

Rolling Back a Failed Upgrade

A Windows 7 upgrade automatically rolls back to Windows Vista if there is a failure during the installation process. You can also roll back to Windows Vista manually up until the point where a successful logon occurs. This means that if there is a problem with a hardware driver that prevents you from successfully logging on, you can go back to your existing Windows Vista installation. After you have performed a successful logon to Windows 7, however, it is not possible to return to Windows Vista without performing a clean installation of the operating system or a restore from backup.

Quick Check

- Under which conditions is it not possible to upgrade from Windows Vista to Windows 7?

Quick Check Answer

- You cannot upgrade from an x86 version of Windows Vista to an x64 version of Windows 7, nor from an x64 version of Windows Vista to an x86 version of Windows 7. It is also not possible to upgrade from certain editions of Windows Vista to certain editions of Windows 7.

Lesson 3: Managing User Profiles

Unless you are performing a direct upgrade from Windows Vista to Windows 7 or one that uses roaming user profiles, any Windows 7 deployment requires that you have a plan for moving user profile data from the user's previous computer to the new computer. Getting user data such as e-mail messages and Web browser bookmarks properly transferred is as important to take into account when performing a Windows 7 deployment as getting the right hardware platform on which to run the operating system. If you cannot get all the users' data that was on their old computers on to their new computers, they may not be able to do their jobs. Users also feel less intimidated by a new operating system if all their old operating system preferences are in effect from the moment they first log on. The more comfortable users are with a new operating system, the more favorably they will look on the transition. In this lesson, you learn how to migrate user data from previous versions of Windows, or from an existing installation of Windows 7, to a new installation of Windows 7.

After this lesson, you will be able to:

- Migrate user profiles from one computer running Windows 7 to another.
- Migrate user profiles from previous versions of Windows.

Estimated lesson time: 40 minutes

Migrating User Profile Data

User data includes more than just documents from a word processor. User data includes things such as favorite Internet sites, customized application settings such as e-mail account data, desktop backgrounds, files, and folders. Unless you are using roaming user profiles in your organization, the computer that a person uses is likely to host important data. Migrating a user from Windows XP or Windows Vista to Windows 7 successfully involves ensuring that all this important data makes the transition from a person's old computer to the new one.

You can view the list of user profiles stored on a computer running Windows 7 by opening System within Control Panel, clicking Advanced System Settings, and then clicking the Settings button in the User Profiles area of the Advanced System Settings tab. From this dialog box, shown in Figure 1-17, you can view the size of user profiles, delete user profiles stored on the computer, or change the user profile from a local profile to a roaming user profile.

A roaming user profile is a profile stored on a server that is accessible from any computer running Windows 7 on a network. Administrators implement roaming user profiles when people do not use a specific computer, but they might log on to any computer in the organization. Roaming user profiles also allow central backup of user data.

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. These scenarios set up real-world situations involving the topics of this chapter and ask you to create a solution.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- Windows 7 requires a 1 GHz x86 or x64 processor, 512 MB of RAM for the Starter or Home Basic editions, and 1 GB of RAM for other editions. To support the Aero UI, Windows 7 requires a graphics card with 128 MB RAM.
- Windows 7 can be dual-booted with Windows XP and Windows Vista. A computer requires multiple partitions to support dual-boot.
- Windows 7 can be installed from DVD-ROM, USB storage device, network share, or WDS.
- The Upgrade Advisor can inform you of which hardware and software attached to your computer running Windows Vista is compatible with Windows 7.
- You can only upgrade from an x86 version of Windows Vista to an x86 version of Windows 7. You can only upgrade from an x64 version of Windows Vista to an x64 version of Windows 7.
- You cannot upgrade directly from Windows XP to Windows 7. You can only upgrade from Windows Vista to Windows 7.
- Windows Easy Transfer allows migration of user profile data from computers running Windows XP and Windows Vista to computers running Windows 7. It is suitable for when a small number of computers need to be migrated.
- USMT is a command-line tool that allows migrating of user profile data from computers running Windows XP and Windows Vista to computers running Windows 7. It is suitable when a large number of computers need to be migrated.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- **dual-boot**
- **Netbook**

4. Save the configuration file to the Iso subdirectory of the Windows PE build directory. The ImageX tool will recognize a Wimscript.ini file in the same location.



EXAM TIP

No command exists that instructs ImageX to detect a Wimscript.ini file. The ImageX tool automatically detects Wimscript.ini if it is saved to the same folder as the ImageX tool.

5. Create an image (.iso) file by using the Oscdimg tool. For example, on an x86 computer you would click All Programs, Microsoft Windows AIK, open the Deployment Tools Command Prompt, and enter the following:

```
oscdimg -n -bc:\winpe_x86\etfsboot.com c:\winpe_x86\ISO  
c:\winpe_x86\winpe_x86.iso
```

The code for an x64 WinPE disk would be:

```
oscdimg -n -bc:\winpe_amd64\ISO\boot\etfsboot.com c:\winpe_amd64\ISO  
c:\winpe_amd64\winpe_amd64.iso
```

MORE INFO ETFSBOOT.COM

This specifies the location of the El Torito boot sector file. For more information, see <http://technet.microsoft.com/en-us/library/cc749036.aspx>. Note also there is no space between the `-b` flag and `C:\Winpe_x86\Etfsboot.com`.

6. Burn the image (Winpe_x86.iso) to a CD-ROM or DVD-ROM disk. Windows AIK does not include CD/DVD-ROM burning software. Use trusted third-party software to burn the image to optical media. You now have a bootable Windows PE optical disk containing the ImageX tool.

Capturing the Installation onto a Network Share

You can capture an image of your reference computer by using Windows PE and the ImageX tool. Then you store that image on a network share. Alternatively, on a computer running Windows 7 Enterprise or Ultimate edition, you can store the image on a VHD and make that VHD bootable, as described in the practice in Lesson 2, later in this chapter.

To capture the installation image you have created on your reference computer to a network share, perform the following procedure:

1. Insert your Windows PE media into your reference computer and restart the computer. As before, you may have to override the boot order to boot from the CD/DVD-ROM drive. If so, select the appropriate function key to override the boot order during initial boot.
2. Windows PE starts and opens a command-prompt window. Use the ImageX tool located on your Windows PE media to capture an image of your reference computer installation. For example, if your optical drive is drive E:, your installation is on drive C:, and you want to capture the image on drive D:, you would enter:

```
e:\imagex.exe /capture C: d:\installationimage.wim "my Win7 Install" /compress  
fast /verify
```

13. Enter **format fs=ntfs label="Windows"**.
14. Enter **assign letter=d**.
15. Enter **exit**.

NOTE CREATING A SCRIPT

You can create a script with this information in a text file and store in the same location as your image. To run the script from a Windows PE command-prompt window, enter `diskpart /s <scriptname>.txt`, where `<scriptname>` is the name of the text file that includes the Diskpart commands. Figure 2-8 shows a typical script file named `DiskConfigurationFormat.txt`.




FIGURE 2-8 A disk configuration format file

16. Copy the image from the network share to your local hard drive. For example, at an elevated command prompt, type:

```
net use y: \\network_share\images
copy y:\installationimage.wim d:
```
17. If necessary, provide network credentials for appropriate access.
18. Apply the image to the hard drive by using the ImageX tool located on the Windows PE media. For example, at an elevated command prompt, enter:

```
e:\imagex.exe /apply d:\myimage.wim 1 C:
```
19. Use BCDboot to initialize the Boot Configuration Data (BCD) store and copy boot environment files to the system partition. For example, at a command prompt, type:

```
d:\windows\system32\bcdboot d:\windows
```

MORE INFO BCDBoot

For more information about BCDBoot, see <http://technet.microsoft.com/en-us/library/cc721886.aspx>.

Your custom image is now deployed onto your destination computer. The computer is ready for customer delivery. Repeat this procedure for each additional computer that you configure.

MORE INFO WINDOWS 7 DEPLOYMENT

For more information about installing and preparing a reference computer, including the use of the Windows SIM tool to prepare answer files for automatic installation, see <http://technet.microsoft.com/en-us/library/dd349348.aspx>.

NOTE CROSS-ARCHITECTURE TOOLS

Both ImageX(x86) and Windows PE(x86) are cross-architecture tools. You can capture both 32-bit and 64-bit images using these utilities.

**EXAM TIP**

Know the functions of a Wimscript.ini, disk configuration format, and Autounattend.xml file. Know how these files are created, how they are accessed, and when you would use them.

Windows Imaging Format

You can use the ImageX Windows AIK tool to create a WIM file that images a reference computer. Unlike ISO files, which are used to contain images of operating systems and toolkits across an intranet or the Internet, WIM is a file-based disk image format that contains a set of files and associated file system metadata. However, unlike sector-based formats (such as ISO) used for CD-ROM and DVD-ROM images, WIM is file-based, which means that the smallest unit of information in a WIM image is a file. A file-based image is hardware-independent and provides unique single-instance storage of a file that can be referenced multiple times in the file system tree.

The files are stored inside a single WIM database. The resource cost of reading or writing many thousands of individual files on a local disk is reduced by hardware- and software-based disk caching and sequential data reads and writes. WIM images are deployed to an existing volume or partition because the toolset does not create low-level disk structures, nor does it format them. Instead, the Microsoft command-line tool Diskpart is used to create and format volumes on the target computer.

TABLE 2-4 DISM Command Options

OPTION	DESCRIPTION	FLAGS
<i>/mount-wim</i>	<p>Mounts the WIM file to the specified directory so that it is available for servicing. The optional <i>/readonly</i> flag sets the mounted image with read-only permissions.</p> <p>Example: <code>dism /mount-wim /wimfile:C:\practice\myimages\install.wim /index:1 /mountdir:C:\practice\offline /readonly</code></p>	<p><i>/wimfile:<path_to_image.wim></i></p> <p><i>/index:<image_index></i></p> <p><i>/name:<image_name></i></p> <p><i>/mountdir:<path_to_mount_directory></i></p> <p><i>/readonly</i></p>
<i>/commit-wim</i>	<p>Applies the changes you have made to the mounted image. The image remains mounted until the <i>/unmount-wim</i> option is used.</p> <p>Example: <code>dism /commit-wim /mountdir:C:\practice\offline</code></p>	<p><i>/mountdir:<path_to_mount_directory></i></p>
<i>/unmount-wim</i>	<p>Dismounts the WIM file and either commits or discards the changes that were made while the image was mounted.</p> <p>Example: <code>dism /unmount-wim /mountdir:C:\practice\offline /commit</code></p>	<p><i>/mountdir:<path_to_mount_directory>{/commit /discard}</i></p>
<i>/remount-wim</i>	<p>Recovers an orphaned WIM mount directory.</p> <p>Example: <code>dism /remount-wim /mountdir:C:\practice\offline</code></p>	<p><i>/mountdir:<path_to_mount_directory></i></p>
<i>/cleanup-wim</i>	<p>Deletes all the resources associated with a mounted WIM image that has been abandoned. This command does not dismount currently mounted images, nor does it delete images that can be remounted.</p> <p>Example: <code>dism /cleanup-wim</code></p>	<p>None</p>
<i>/get-wiminfo</i>	<p>Displays information about the images within the WIM. When used with the <i>/index</i> option, information about the specified image is displayed.</p> <p>Example: <code>dism /get-wimInfo /wimfile:C:\practice\offline\install.wim /index:1</code></p>	<p><i>/wimfile:<path_to_image.wim></i></p> <p><i>/index:<Image_index></i></p> <p><i>/name:<Image_name></i></p>

TABLE 2-7 Sysprep Log File Locations

ITEM	LOG PATH
<i>Generalize pass</i>	%WINDIR%\System32\Sysprep\Panther
<i>Specialize pass</i>	%WINDIR%\Panther\
Unattended Windows setup actions	%WINDIR%\Panther\Unattendgc

PRACTICE Creating a WIM Image

In this practice, you install the Windows AIK. You then create a Windows PE boot disk and boot the computer into Windows PE. This enables you to use the ImageX tool in the Windows AIK to create a WIM image of the computer.

EXERCISE 1 Installing the Windows AIK and Creating a Windows PE Boot DVD

In this exercise, you download the ISO image in Windows AIK and create an installation DVD. You then install the Windows AIK. Instructions for doing this were given in the section entitled “Installing and Using the Windows Automated Installation Toolkit,” earlier in this lesson. You create a Windows PE build directory and copy ImageX into it. You use the Oscdimg tool to create an ISO image of Windows PE. You burn this image onto optical media (CD-ROM or DVD) that you can use to boot the computer. You need to be connected to the Internet to perform this exercise.

1. Log on to the Canberra computer using the Kim_Akers account.
2. Download the appropriate ISO image, burn this to optical media, and install the Windows AIK.
3. In Accessories in the All Programs menu, right-click Command Prompt and choose Run As Administrator. If prompted, click Yes to permit the program to run.
4. In the Command Prompt window, enter **cd C:\Program Files\Windows AIK\Tools\PETools**.
5. At the C:\Program Files\Windows AIK\Tools\PETools> prompt, enter **copy.cmd x86 c:\winpe_x86**. This exercise is written for a 32-bit computer and the Windows PE build directory is Winpe_x86. If you are using an amd64 or ia64 computer, amend the entry accordingly. Figure 2-10 shows the output from this command. Next, copy winpe.wim to the ISO\Sources folder and rename it boot.wim: **xcopy /chery c:\winpe_x86\winpe.wim c:\winpe_x86\ISO\Sources\boot.wim**.
6. To copy ImageX into the Windows PE build directory, enter **copy "c:\program files\Windows AIK\Tools\x86\imagex.exe" c:\winpe_x86\iso**.
7. To create an image (.iso) file by using the Oscdimg tool, click Microsoft Windows AIK in All Programs and then click Deployment Tools Command Prompt.

9. The ISO image is in C:\Winpe_x86 and is named Winpe_x86.iso.

EXERCISE 2 Creating a WIM Image of the Canberra Computer

In this exercise, you boot the Canberra computer from the optical Windows boot disk that contains ImageX, which you created in Exercise 1. You then create a WIM image of the Windows 7 installation and (optionally) save it to a network share.

1. If necessary, log on to the Canberra computer using the Kim_Akers account.
2. On the Canberra computer, insert the Windows PE medium and restart the computer.

NOTE CHANGING THE BIOS BOOT ORDER

To boot from the optical drive, you may have to override the BIOS boot order. During initial boot, select the appropriate function key.

3. Windows PE starts and opens a command-prompt window.
4. To capture an image of the reference installation by using the ImageX tool located on your Windows PE medium, enter **e:\imagex.exe /capture c: d:\images\myimage.wim "Canberra Win7 Install" /compress fast /verify**. This command uses ImageX on the CD/DVD-ROM drive E: to capture the image of the system disk C: to the folder images on the second hard disk D:. If your volume assignments are different, amend the command accordingly. The command takes a considerable time to complete and lists folders (such as the recycle bin) that are not included in the image by default.
5. Enter **exit** and remove your Windows PE boot disk. The computer boots into Windows 7.
6. Check that the file Myimage.wim exists on the D: drive (or wherever you chose to put it).
7. Optionally, if you want to share the image across a network, create a network share, (for example, \\Canberra\Images) and map it to a network drive (such as Y:) and then copy the WIM file to this share.

Lesson Summary

- The Windows AIK introduced in Windows 7 offers various tools for creating system images. These include Windows SIM, ImageX, Oscdimg, DISM, USMT, and several Windows PE tools.
- You use Windows SIM to create an unattend answer file that you can in turn use with a WIM image to install a reference computer. You use Sysprep to prepare the image and then boot the reference computer into Windows PE and use the ImageX tool to capture the image in a WIM file.

Disk dialog box. You select the partition and click OK. Typically, you do not need to change the default settings. The status of the disk then changes to Online.

You create a new simple volume on a VHD by right-clicking Unallocated and selecting New Simple Volume. This starts the New Simple Volume Wizard. You specify size, file system, and drive letter; label the drive; and click Finish to create the VHD.

Attaching and Detaching a VHD

You can also use the Disk Management tool to attach a VHD so you can use it and to detach it so you can change its properties or delete it. In Computer Management, you click Disk Management and then right-click Disk Management and click Attach VHD. This opens the Attach Virtual Hard Disk dialog box. Click OK to attach the existing VHD. If you do not want to change the VHD contents (for example, if you have installed an operating system on it), you can select the Read-Only check box.

To detach a VHD, you click the icon beside the disk designation and click Detach VHD. A Detach Virtual Hard Disk message appears. Click on OK to detach the VHD. If you want to delete the VHD permanently after it is detached, you can select the Delete The Virtual Hard Disk File After Deleting The Disk check box.

Using the Diskpart Utility to Create and Attach a VHD

You can use the Diskpart command-line utility to create and attach a VHD by performing the following steps:

1. On the Accessories menu, right-click Command Prompt and choose Run As Administrator. If necessary, click Yes to allow the program to run.
2. Enter **diskpart**.
3. Enter **create vdisk file=c:\win7\myothervhd.vhd maximum=20000**. This creates a VHD file called Myothervhd Win7 with a maximum size of 20 GB in a folder called Win7 on the C: drive. You can also create a VHD on a second internal hard disk or on a USB external hard disk formatted with the NTFS filing system.
4. Enter **select vdisk file=c:\win7\myothervhd.vhd**.
5. Enter **attach vdisk**.
6. Enter **create partition primary**.
7. Enter **assign letter=v**.
8. Enter **format quick label=Windows7**.
9. Enter **exit**.

This creates the VHD file *C:\Win7\Myothervhd.vhd* as a primary partition. Figure 2-12 shows the Diskpart commands to create and attach a new VHD. Figure 2-13 shows the newly attached disk in Disk Management with drive letter V:.

This deletes the specified operating system entry from the store and removes the entry from the display order.

When you restart your computer after successfully completing this procedure, you should see an additional entry in the Boot menu along with the default Windows 7 operating system.

MORE INFO BCDEdit

For more information about BCDEdit, go to <http://msdn.microsoft.com/en-us/library/aa906217.aspx>, expand BCD Boot Options Reference, and click the links in the navigation pane.



EXAM TIP

You can use Bcdedit.exe to enable a VHD file as a boot option, but you cannot use the tool to create VHD files.

Using the Windows Image to Virtual Hard Disk Tool

You can use the WIM2VHD command-line tool to create VHD images from any Windows 7 installation source or from an image in a custom WIM file. WIM2VHD creates VHDs that boot directly to the Out-of-Box Experience (OOBE). You can also automate the OOBE configuration by supplying your own Unattend.xml file.

You need a client computer running Windows 7 that has the Windows AIK installed, and an operating system image in a WIM file. You also need to have created a native VHD on that computer.

The WIM2VHD tool runs from the Cscript command. The syntax is as follows:

```
cscript wim2vhd.wsf /wim:<wimPath> /sku:<sku> [/vhd:<vhdPath>] [/size:<vhdSizeInMb>]
[/disktype:<dynamic|fixed>] [/unattend:<unattendxmlPath>] [/qfe:<qfe1,..,qfen>]
[/hyperv:<>true|false>] [/ref:<ref1,..,refn>] [/dbg:<args>] [/passthru:<physicaldrive>]
```

WIM2VHD Parameters

Table 2-8 describes the parameters of the WIM2VHD tool.

TABLE 2-8 WIM2VHD Parameters

PARAMETER	DESCRIPTION
/wim:<wimPath>	This is the path of the WIM file you use when creating the VHD.
/sku:<skuName> <skuIndex>	The Stock-Keeping Unit (SKU) identifies the operating system to use when creating the VHD (for example, "HomePremium"). You can also specify a number that you obtain by using ImageX to analyze the relevant WIM file.

PARAMETER	DESCRIPTION
<i>/vhd:</i> <vhdPath> (optional)	This defines the path and the name of the VHD to be created. If a file with this name already exists, it will be overwritten. If no VHD is specified, a VHD will be created in the current folder.
<i>/size:</i> <vhdSizeInMb> (optional)	For Fixed disks, this is the size in megabytes of the VHD that will be created. For Dynamic disks, this is the maximum size in megabytes to which the VHD can grow if additional space is required. If you do not specify this parameter, a default value of 40 GB is used.
<i>/disktype:</i> <dynamic fixed> (optional)	This specifies what kind of VHD should be created, dynamic or fixed. A Fixed disk allocates all of the necessary disk space for the VHD upon creation. A Dynamic disk only allocates the space required by files in the VHD at any given time and will grow as more space is required. The default value is Dynamic.
<i>/unattend:</i> <unattendxmlpath> (optional)	This specifies the path to an Unattend.xml file that is used to automate the OOBE portion of Windows setup the first time the VHD is booted.
<i>/qfe:</i> <qfe1,..,qfen> (optional)	This is a comma-separated list of Quick Fix Engineering (QFE) or hotfix patches to apply to the VHD after the WIM is implemented.
<i>/ref:</i> <ref1,..,refn> (optional)	This is a comma-separated list of WIM "pieces" (split files) to apply to the VHD. A WIM "piece" is the result of a split WIM, and typically has a .swm file extension. The first piece of the split WIM should be specified with the <i>/wim</i> switch. Subsequent pieces should be specified (in order) with <i>/ref</i> .
<i>/dbg:</i> <protocol>,<port/channel/target>[,<baudrate>]. (optional)	This configures debugging in the OS on the VHD.

You can use your own custom WIM files in this process. However, be careful. Although Microsoft supports the underlying process, as documented in the Windows AIK, WIM2VHD is not supported at this time.

You can copy files manually into the VHD, but there is no mechanism to do this with WIM2VHD.

WIM2VHD Examples

To create a Windows 7 Ultimate VHD with an automated setup answer file Unattend.xml, open an elevated command prompt and enter:

```
cscript wim2vhd.wsf /wim:x:\mysources\install.wim /sku:ultimate /unattend:C:\answer_files\unattend.xml
```

You need to adjust the location of the WIM file and the answer file to your own specifications.

To apply the first image in a custom WIM in the folder C:\Mystuff to a VHD named Mycustom.vhd when you have analyzed the WIM file with ImageX and know the SKU is designated as 1 within the WIM, open an elevated command prompt and enter:

```
cscript wim2vhd.wsf /wim:C:\mystuff\custom.wim /sku:1 /VHD:C:\mycustom.vhd
```

MORE INFO WIM2VHD

For more information about WIM2VHD, see <http://code.msdn.microsoft.com/wim2vhd>.

Using the Offline Virtual Machine Servicing Tool to Update a VHD

The Offline Virtual Machine Servicing Tool 2.0.1 is a solution accelerator (as is MDT 2010). In addition to the appropriate installation files, a *solution accelerator* provides automated tools and additional guidance files. You can install the tool on a server running Windows Server 2008 or Windows Server 2003 SP2, where it works with Microsoft System Center Virtual Machine Manager (SCVMM) 2007 or SCVMM 2008 to maintain offline virtual machines and VHDs.

If your server is on the same network as a client running Windows 7 Enterprise or Ultimate edition, on which you have configured a bootable VHD, you can use the tool to update the VHD content when the VHD is typically offline. If your computer running Windows 7 is not normally booted from the VHD, the offline VHD does not receive operating system updates. The tool provides a way to keep the VHD up to date so that booting from the VHD does not introduce vulnerabilities into your computer.

The Offline Virtual Machine Servicing Tool can be configured to boot the client computer from the VHD just long enough for the VHD to receive updates from either SCCM 2007 or Windows Server Update Services (WSUS). As soon as the VHD's operating system is up to date, the tool reboots the client computer from its default boot disk.

The Offline Virtual Machine Servicing Tool solution accelerator includes the following features:

- Brief Overview
- OfflineVMServicing_x64.msi installation file
- OfflineVMServicing_x86.msi installation file
- Offline Virtual Machine Servicing Tool Getting Started Guide
- Offline_VM_Servicing_Tool_2.0_Release_Notes
- Offline_Virtual_Machine_Servicing_Tool_Help

The tool uses a servicing job that you schedule using Windows Task Scheduler on the server to manage the update operation. The servicing job boots the client computer from the

PRACTICE Creating a VHD

In this practice, you use the Computer Management tool to create a VHD. You need to have completed the practice exercises in Lesson 1 before attempting this practice.

EXERCISE 1 Creating a VHD

To use Computer Management to create a VHD, perform the following procedure:

1. Log on to the Canberra computer with the Kim_Akers account.
2. Create a folder called VHDs on the C: drive. If you prefer to use an external USB disk drive, adjust your drive letter accordingly, but first ensure that the external drive is formatted with the NTFS file system.
3. On the Start menu, right-click Computer and choose Manage. If prompted, click Yes to allow the program to run.
4. Select Disk Management.
5. Right-click Disk Management and choose Create VHD, as shown in Figure 2-15.




FIGURE 2-15 Creating a VHD

8. Ensure MBR (Master Boot Record) is selected and click OK. The status of the disk changes to Online.
9. On the newly created disk, right-click Unallocated and select New Simple Volume. This starts the New Simple Volume Wizard. Click Next.
10. Click Next to accept the volume size defaults.
11. In the Assign Drive Letter Or Path dialog box, select W and then click Next.
12. In the Format Partition dialog box, give the volume a label (such as MyVHD), as shown in Figure 2-18. Ensure that Perform A Quick Format is selected. Click Next.




FIGURE 2-18 The Format Partition dialog box

13. Click Finish.

Before You Begin

To complete the exercises in the practices in this chapter you need to have done the following:

- Installed the Windows 7 operating system on a stand-alone client PC, as described in Chapter 1, “Install, Migrate, or Upgrade to Windows 7.” You need Internet access to complete the exercises.
- Completed all the practice exercises in Chapter 2, “Configuring System Images.” In particular, you need to have installed the Windows Automated Installation Kit (Windows AIK) and deployed an offline image of the Canberra computer on a bootable virtual hard disk (VHD).



REAL WORLD

Ian McLean

Each of the tools you use for network and system administration has its own set of features and enables you to perform specific tasks. Where people sometimes get upset and confused is if there is overlap. For example, you can use Windows Deployment Services (WDS) or the Microsoft Deployment Toolkit (MDT) 2010 to deploy Windows 7 images to client computers. However, MDT 2010 allows you to specify a set of configuration tasks that should be run on a computer after an image has been deployed to it through WDS, whereas WDS cannot run configuration tasks on a client after the image has been deployed.

You use ImageX to create an image of a computer running Windows 7 while it is booted to Windows Preinstallation Environment (Windows PE) and you use the Deployment Image Servicing and Management (DISM) tool to manipulate that image after it has been created. However, you can use ImageX or DISM to mount an image so you can work with it offline.

So, you are entitled to ask, which tool do I use and, more to the point in this book, what tool will the examination ask about? The simple answer is to use the most recently introduced tool when there is a choice. For example, ImageX has been around for some time, whereas DISM was introduced fairly recently; however, ImageX has new features in the latest edition of the Windows AIK.

Traditionally, examinations ask more questions about new features of an operating system and new tools that are introduced to carry out tasks on the new operating system than they do about features that are unchanged from the previous operating system and tools that, however worthy, have been around for a while. This is simply a statement of fact and any conclusions you draw from it are your own.

Lesson 1: Managing a System Image Before Deployment

Sometimes when you have created a master reference image for deploying to other computers, you might find that you need to amend it. You might need to add a new driver, change settings, and support multiple languages. Even a fairly minor change, such as enabling a feature currently disabled in the image, can generate a considerable workload if it needs to be done after the image has been distributed to several hundred computers.

Typically, it involves less administrative effort if you make these changes without deploying the image and recapturing it. If your requirement is to add security updates, then it is certainly preferable to apply the security patches offline—otherwise, you are deploying an insecure image. If you service the image offline, you do not need to run the Sysprep tool and therefore do not need to use a rearm parameter. Finally, you might want to apply an amended Autounattend.xml file for unattended install or an additional Unattend.xml file that automates post-installation tasks such as installing mission-critical applications. Unattend.xml files are discussed in the section entitled “Unattended Servicing Command-Line Options,” later in this chapter.

This lesson discusses how you use Windows AIK tools such as ImageX and DISM to mount a system image and how you use DISM to manage the image, insert packages, insert updates, enable and disable features, manage international settings, manage language packs, and associate unattend answer files.

After this lesson, you will be able to:

- Mount an offline image for servicing.
- Use DISM to manage and manipulate the image.
- Associate one or more answer files with the image.
- Commit and unmount the image.

Estimated lesson time: 50 minutes

Using DISM WIM Commands and Mounting an Image

Chapter 2 introduced the Windows 7 DISM command-line tool. DISM enables you to service offline images, mount and dismount Windows Imaging format (WIM) files, and customize Windows PE boot images. The DISM tool replaces many of the tools in previous versions of the Windows AIK, including Package Manager (Pkgmgr.exe), the International Settings Configuration Tool (Intlcfg.exe), and the Windows PE command-line tool (PEimg.exe).

Microsoft has designed the DISM tool to manage WIM images. Also, DISM is backward-compatible with Vista tools, such as Pkgmgr.exe, Intlcfg.exe, and PEimg.exe, so scripts that you developed and tested for Vista should work unamended in Windows 7.

Any parameter that can be used online can also be used offline by specifying a mounted WIM image with the */image* switch. For example, the following command lists all the drivers in the image mounted in the folder C:\MountedImages:

```
dism /image:c:\mountedimages /get-drivers /all
```

Table 3-2 lists the information retrieval parameters that you can use with an offline mounted image but not with an online image.

TABLE 3-2 Parameters That Cannot Be Used with an Online Image

PARAMETER	DESCRIPTION
/Get-AppPatchInfo	Displays information about installed Windows Installer patch files (MSP patches)
/Get-AppPatches	Displays information about all applied MSP patches for all installed applications
/Get-AppInfo	Displays information about a specific installed Windows Installer (MSI) application
/Get-Apps	Displays information about all installed MSI applications

Servicing Drivers, Applications, Patches, Packages, and Features

You can use driver servicing commands on an offline mounted image to add and remove drivers based on the .inf file format. You can specify a directory where the driver .inf files are located, or you can point to a driver by specifying the name of the .inf file.

On an online running operating system, you can only enumerate drivers and obtain driver details. The commands and options to list drivers and obtain driver information were discussed in the previous section of this lesson. DISM can manage only .inf file drivers. Windows Installer (MSI) and other driver package types (such as .exe files) are not supported.

The following driver servicing options are available for an offline image:

```
dism /image:path_to_image_directory [/get-drivers | /get-driverinfo | /add-driver | /remove-driver]
```

For example, if you wanted to add the driver Mydriver.inf that you have downloaded and stored in the folder C:\Newdrivers, you would use a command similar to the following:

```
dism /image:c:\mountedimages /add-driver:c:\newdrivers\mydriver.inf
```

Figure 3-8 shows the output from this command.

On x64-based computers running Windows 7, drivers must have a digital signature. However, you might want to install an unsigned driver for test purposes. In this case you can use the */forceunsigned* parameter to override this requirement.

You can use the */remove-drivers* option to remove third-party drivers from an offline image. You cannot remove default drivers with the *dism* command. When you add third-party drivers, they are named *oem0.inf*, *oem1.inf*, and so on. You must specify the published name, but fortunately, the */get-drivers* parameter lists both the published name and the original name. If you have installed a lot of third-party drivers and are having difficulty finding the new name of the driver you want to remove, you can direct the output from a *dism* command that uses the */get-drivers* parameter to a text file and search this file for the original name. When you have identified the driver's published name, such as *oem10.inf*, you can then remove it using a command similar to the following:

```
dism /image:c:\mountedimages /remove-driver /driver:oem10.inf
```

Servicing Applications and Application Patches

You can use application servicing command-line options applied to a offline image to check the applicability of Windows Installer application patches and to query the offline image for information about installed Windows Installer applications (.msi files) and application patches (.msp files).

None of the application servicing commands can be applied to online images, although if an image is online, it can receive updates from, for example, Windows Server Update Services (WSUS) or Microsoft Update. If you are administering an enterprise network, you should consider the Offline Virtual Machine Servicing Tool and System Center Virtual Machine Manager, which were discussed in Chapter 2.

DISM offers the following options available to list Windows Installer (.msi) applications and .msp application patches, and to check the applicability of an application patch on an offline system image:

```
dism /image:path_to_directory [/check-apppatch | /get-apppatchinfo: | /get-apppatches | /get-appinfo | /get-apps]
```

When managing applications and patches, your first step should be to discover what application patches and applications exist and are applicable to the image. For example, in an image mounted directly from an Install.wim file copied from the installation media, it is likely that no applicable patches or applications exist. To obtain information about application patches (MSP files) applicable to a mounted image, you would use a command similar to the following:

```
dism /image:c:\mountedimages /get-apppatches
```

If you know the product code globally unique identifier (GUID) of a Windows Installer application, you can use the */productcode* parameter to display all the application patches in the specified application. You would use a command similar to the following:

```
dism /image:c:\mountedimages /get-apppatches /productcode:{GUID}
```

If you want to display information about specific .msp patches applicable to the offline image, you can use the `/check-apppatch` parameter. You use `/patchlocation` to specify the path to the MSP patch file. You can specify multiple patch files by using `/patchlocation` more than once in the command. For example, to display information about two patch files, 30880d0.msp and 8c82a.msp (both in C:\Windows\Installer) in the mounted image, you would enter the following command:

```
dism /image:c:\mountedimages /check-apppatch
/patchlocation:c:\windows\installer\30880d0.msp
/patchlocation:c:\windows\installer\8c82a.msp
```

If you need detailed information about all installed MSP patches applicable to the offline image, you would enter a command similar to the following:

```
dism /image:c:\mountedimages /get-apppatchinfo
```

✓ Quick Check

- You want to add all the drivers in the folder C:\Orinsnewdrivers and its subfolders to the mounted offline image in D:\Orinsimage. What command would you use?

Quick Check Answer

- `dism /image:d:\orinsimage /add-driver /driver:c:\orinsnewdrivers /recurse`

You can use the `/get-apppatches` option described earlier in this section to find the patch code GUID and the product code GUID specific to a patch. You can also use the `/get-apps` option described here to list all product code GUIDs for an installed Windows Installer application. You can filter the information returned by the `/get-apppatchinfo` parameter either by the patch code GUID or the product code GUID, or by both, for example:

```
dism /image:c:\mountedimages /get-apppatchinfo /patchcode:{patch_code_GUID}
/productcode:{product_code_GUID}
```

In addition to obtaining information about applicable application patches, you typically need to obtain information about the MSI applications. The `/get-apps` parameter lists the MSI applications installed on the mounted image and you can use it to determine each application's GUID; for example:

```
dism /image:c:\mountedimages /get-apps
```

You can then obtain more detailed information about installed applications by using the `/get-appinfo` parameter. Optionally you can filter this information by specifying the product code GUID for the application in which you are interested; for example:

```
dism /image:c:\mountedimages /get-appinfo /ProductCode:{product_code_GUID}
```

If you do not specify a product code GUID, the `/get-appinfo` parameter returns detailed information about all installed MSI applications.



EXAM TIP

Remember that */get-apppatches* and */get-apppatchInfo* apply only to installed patches (.msp files) and that */get-apps* and */get-appinfo* apply only to Windows Installer applications (.msi files). You cannot, for example, use DISM to obtain information about .exe or .dll files. Also, remember that when you check the applicability of an MSP patch, only the Windows Installer applications for which the patch is applicable will be displayed. One patch can be applied to many installed applications and many patches can be applied to one application.

Adding Applications to an Image

The */check-apppatch*, */get-apppatchinfo*, */get-apppatches*, */get-appinfo*, and */get-apps* DISM options obtain information about Windows Installer applications and installed patches on an offline mounted image. The next section describes how you add cabinet (.cab) or Windows Update Stand-alone Installer (.msu) files to an image and, in particular, install security patches to offline-mounted images. You can also enable and disable Windows features, but you cannot add features or any other type of executable files, such as .exe, .bat, .com, or .vbs files. The DISM command does not have an */add-apps* option.

If you want to add a mission-critical application to the image for distribution, you can install that application on your reference computer before you image it. If, however, you want to add an application to an already existing offline image, the DISM tool does not do this. Instead, you should use the Add Application Wizard provided by MDT 2010, which is discussed in Lesson 2, “Deploying System Images.”

You can also use DISM to associate an image with an Unattend.xml answer file. Such a file automates installation of the image but also automates post-installation tasks, for example, connecting to a file server and installing applications or configuring settings. This approach, where applications and settings are applied after installation rather than included in the image, is known as “thin image” and is described in Lesson 2 of this chapter.

Servicing Operating System Packages

One of the problems you have with system images either held for distribution to a number of computers or installed on a bootable VHD of a single client computer for failover purposes is that you need to keep the image up to date, particularly with security updates and fixes. Otherwise, if you boot with the new image, the computer is vulnerable to known security threats.

Package Installation Considerations

When you install a package in an offline image, the package state becomes “install pending,” and the package is installed when the image is booted and pending online actions are processed. If subsequent actions are requested, they cannot be processed until the previous pending online actions complete. If a package is in the “install pending” state and you stage the package, the package state is set to “uninstall pending” because the package must be uninstalled before it can be staged.

Some packages require other packages to be installed first. If there are dependency requirements, you should use an answer file to install the necessary packages. By passing an answer file to DISM, you can install multiple packages in the correct order. Microsoft recommends the use of an answer file for installing multiple packages. Packages are installed in the order that they are listed in the command line, which in turn can be generated in an answer file.

When you use DISM to list the feature packages in a Windows PE image, the packages will always be listed as pending even when the servicing operation was successful. This is by design, and you do not need to take any further action.

Configuring International Settings in an Image

You can use the DISM tool to manage international settings in a Windows 7 (or a Windows PE) image. You can also query existing settings in an offline or an online image.


You can use the following international servicing options on an offline-mounted image:

```
dism /image:path_to_offline_image_directory [/get-intl] [/set-uilang |  
/set-uilangfallback | /set-syslocale | /Set-UserLocale | /Set-InputLocale | /Set-AllIntl  
| /Set-Timezone | /Set-SKUIntlDefaults | /Set-LayeredDriver] [/Gen-Langini |  
/Set-SetupUILang | /Distribution]
```

You can use the following command on a running operating system:

```
dism /online /get-intl
```

This is the only international servicing option you can apply to a running operating system. Its output is shown in Figure 3-11.



```
Administrator: Deployment Tools Command Prompt
C:\Program Files\Windows AIK\Tools\PETools>dism /online /get-intl
Deployment Image Servicing and Management tool
Version: 6.1.7100.0
Image Version: 6.1.7100.0
Reporting online international settings.
Default system UI language : en-US
System locale : en-US
Default time zone : GMT Standard Time
Active keyboard(s) : 0409:00000409
Keyboard layered driver : PC/AT Enhanced Keyboard <101/102-Key>
Installed language(s): en-US
Type : Fully localized language.
The operation completed successfully.
C:\Program Files\Windows AIK\Tools\PETools>_
```

FIGURE 3-11 International settings for an online operating system

You can set the scratch space and the target path by using commands similar to the following:

```
dism /image:c:\mountedpeimage /set-scratchspace:256
dism /image:c:\mountedpeimage /set-targetpath:D:\WinPEboot
```

Scratch space is specified in megabytes. Valid values are 32, 64, 128, 256, and 512. In hard disk boot scenarios, the target path defines the location of the Windows PE image on the disk. The path must be at least 3 characters and no longer than 32 characters. It must have a volume designation (C:\, D:\, and so on) and it must not contain any blank spaces.

File logging (or profiling) lets you create your own profiles in Windows PE 3.0 or later. By default, profiling is disabled. You can enable it, or disable it if previously enabled, by entering commands similar to the following:

```
dism /image:c:\mountedpeimage /enable-profiling
dism /image:c:\mountedpeimage /disable-profiling
```

When you create one or more profiles, each is stored in its own folder and identified in the file Profile.txt. You can remove any files from a Windows PE image that are not part of the custom profiles and check the custom profile against the core profile to ensure that custom application files and boot-critical files are not deleted, by entering a command similar to the following:

```
dism /image:c:\mountedpeimage /apply-profiles:c:\peprofiles\profile01\profile.txt,
c:\peprofiles\profile02\profile.txt
```

The paths to one or more profile.txt files are included in the command as a comma-separated list.

Quick Check

1. You want to obtain a list of PE settings in a mounted Windows PE image in the folder C:\Mypeimage. What command do you enter in the elevated Deployment Tools command prompt?
2. You need to determine the amount of Windows PE system volume scratch space available on a Windows PE system volume in a mounted Windows PE image in the folder C:\Mypeimage when booted in RAMdisk mode. What command do you enter in the elevated Deployment Tools command prompt?

Quick Check Answers

1. `dism /image:c:\mypeimage /get-pesettings`
2. `dism /image:c:\mypeimage /get-scratchspace`

Unattended Servicing Command-Line Options

You can use DISM to apply an Unattend.xml answer file to an image. Typically, you would use this feature when you are installing multiple packages to the image. As stated previously in this lesson, some packages require other packages to be installed first. Microsoft recommends that the best way of ensuring the correct installation order is to use an answer file. If you use DISM to apply an Unattend.xml answer file to an image, the unattended settings in the offline Servicing configuration pass (previously described in Chapter 2) are applied to the Windows image.

The following servicing options are available to apply an Unattend.xml answer file to an offline Windows image:

```
dism /image:path_to_image_directory /apply-unattend:path_to_unattend.xml
```


The following command applies an Unattend.xml answer file to a running operating system:

```
dism /online /apply-unattend:path_to_unattend.xml
```

For example, if the Unattend.xml file is located in C:\Windows\Panther, you can apply it to an offline-mounted image in C:\Mountedimages by entering the following command:

```
dism /image:c:\mountedimages /apply-unattend:c:\windows\panther\unattend.xml
```

Figure 3-14 shows the output from this command. It tells you the answer file has been applied but gives no additional information.



```
Administrator: Deployment Tools Command Prompt
C:\Program Files\Windows AIK\Tools\PETools>dism /image:d:\mountedimages /apply-unattend:c:\windows\panther\unattend.xml
Deployment Image Servicing and Management tool
Version: 6.1.7100.0
Image Version: 6.1.7100.0
The operation completed successfully.
C:\Program Files\Windows AIK\Tools\PETools>
```

FIGURE 3-14 Applying an answer file to an offline-mounted image

Using Answer Files with Windows Images

The ability to associate an Unattend.xml answer file to an image provides a powerful tool to implement and configure image deployment, and to determine actions that can be taken if deployment fails or after deployment succeeds.

When adding data, such as additional drivers or applications, take care that you do not overwrite Windows System files. Overwriting system files can corrupt your computer's operating system.

MORE INFO ADDING APPLICATIONS, DRIVERS, PACKAGES, FILES, AND FOLDERS

For more information about adding applications, drivers, packages, files, and folders, see [http://technet.microsoft.com/en-us/library/dd744568\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd744568(WS.10).aspx).

Using Multiple Answer Files

You can use multiple second answer files (Unattend.xml) to create different custom images. For example, you could create a generic answer file that is used for each of your systems and then apply a second answer file during audit mode for changing disk configurations, drivers, or applications. To do this, you would use the Sysprep command (described in Chapter 2) with the */unattend:answerfile* option. You can run this command manually during audit mode or you can add a custom command.

MORE INFO ADDING CUSTOM COMMANDS AND SCRIPTS

For more information about adding custom commands and scripts, see [http://technet.microsoft.com/en-us/library/dd744393\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd744393(WS.10).aspx).

PRACTICE Mounting an Offline Image and Installing Language Packs


In this practice, you use both ImageX and DISM to mount an image. You also practice unmounting an image. You then apply language packs to a mounted image.

EXERCISE 1 Mounting, Unmounting, and Remounting an Image

In this exercise, you use the ImageX tool to mount the system image Myimage.wim that you installed on the VHD to which you allocated drive letter W:. You mount the image in the folder C:\Mountedimages. You then unmount the folder. Finally, you use DISM to mount the image to the folder D:\Mountedimages. Note that it is not essential to create a different folder to hold the second mounted image. However, if you do not, it is a good idea to delete and re-create the original folder because DISM sometimes returns an error even though the image has been unmounted from the folder. Proceed as follows:

1. Log on to the Canberra computer with the Kim_Akers account.
2. Create a folder called C:\MountedImages. If this folder already exists, ensure that it is empty.
3. On the Start menu, right-click Computer and choose Manage. Select Disk Management. If the W: disk does not appear in the Volume list, right-click Disk

6. If you want, you can unmount the image and then use the DISM tool to mount it in a different folder. To unmount the tool, enter the command: **imagex /unmount c:\mountedimages**. Figure 3-17 shows the output from this command.




```
Administrator: Deployment Tools Command Prompt
C:\Program Files\Windows AIK\Tools\PETools>imagex /unmount c:\mountedimages
ImageX Tool for Windows
Copyright (C) Microsoft Corp. All rights reserved.
Version: 6.1.7100.0
Unmounting: [c:\mountedimages]...
[ 100% ] Mount cleanup progress
Successfully unmounted image.
Total elapsed time: 3 min 44 sec

C:\Program Files\Windows AIK\Tools\PETools>
```

FIGURE 3-17 Unmounting an image

7. Create a folder called D:\MountedImages. If you do not have a second hard disk, you can use the C:\MountedImages folder, but you might need to delete and re-create it if DISM returns an error.
8. Mount the image with the DISM tool by entering the command: **dism /mount-wim /wimfile:w:\myimage.wim /index:1 /mountdir:d:\mountedimages**.
9. Test the image is mounted correctly by entering the command **dism /get-mountedwiminfo**. The output from this command is shown in Figure 3-18. Note that the mounted image folder is not the same as that shown in Figure 3-4 earlier in this lesson.



```
Administrator: Deployment Tools Command Prompt
C:\Program Files\Windows AIK\Tools\PETools>dism /get-mountedwiminfo
Deployment Image Servicing and Management tool
Version: 6.1.7100.0
Mounted images:
Mount Dir : d:\mountedimages
Image File : w:\myimage.wim
Image Index : 1
Mounted Read/Write : Yes
Status : Ok
The operation completed successfully.
C:\Program Files\Windows AIK\Tools\PETools>
```

FIGURE 3-18 An image mounted in D:\Mountedimages

PE media and use ImageX to install the image. Your final step, to make the image bootable, is to use BCDboot from Windows PE to initialize the BCD store and copy boot environment files to the system partition. When you reboot each new computer, it will boot into Windows 7 and will have the same settings configured and applications installed as your original computer. Take care you are not violating any licensing conditions.

PRACTICE Downloading, Installing, and Configuring MDT 2010

In this practice, you download the MDT 2010 installation and documentation files and then install the toolkit. You use the Deployment Workbench tool to create a Distribution Share and install an image.

NOTE THE MDT 2010 INTERFACE

At the time of this writing, MDT 2010 is in beta. Therefore, its eventual interface might vary from what you see in this book.

EXERCISE 1 Downloading the MDT 2010 Installation Files and Documentation

In this exercise, you download the MDT and its associated documentation by accessing <http://www.microsoft.com/downloads/details.aspx?FamilyId=3BD8561F-77AC-4400-A0C1-FE871C461A89&displaylang=en>. You probably first need to supply your Microsoft password credentials. You have the option of downloading the following files:

- MicrosoftDeploymentToolkit_x64.msi
- MicrosoftDeploymentToolkit_x86.msi
- Quick Start Guide for Lite Touch Installation.docx
- Release Notes.docx
- What's New in MDT 2010 Guide.docx

You can download and install the version suitable for your operating system—this book assumes the 32-bit (x86) version. You need no additional software to run the MDT on Windows 7, although if you choose to use the MDT in conjunction with SCCM 2007 on a deployment server, you need to install the relevant software and additional software such as SQL Server.

To download MDT 2010 and its associated documentation, proceed as follows:

1. Log on to the Canberra computer with the Kim_Akers account.
2. Create a folder to hold the downloaded files; for example, C:\Windows 7\MDT 2010 Files. Also create a folder to hold documentation, such as C:\Windows 7\MDT 2010 Documentation.
3. Open Internet Explorer and access <http://www.microsoft.com/downloads/details.aspx?FamilyId=3BD8561F-77AC-4400-A0C1-FE871C461A89&displaylang=en>. If asked, supply your Microsoft Password details.
4. Click Microsoft Deployment Toolkit 2010.
5. Under Microsoft Deployment Toolkit (MDT) 2010, click Download.

However, users cannot be permitted to plug any device they please into their work computers. Device driver software runs as if it is a part of the operating system with unrestricted access to the entire computer, and only authorized device drivers can be permitted.

When a user inserts a device, Windows 7 detects the new hardware and the Plug and Play service identifies the device and searches the driver store for a driver package that matches the device. If a suitable driver is found, the device is considered to be authorized and the Plug and Play service copies the driver file (or files) from the driver store to its operational location, typically C:\Windows\System32\Drivers. The Plug and Play Service configures the registry and starts the newly installed driver.

Installing Device Drivers from Windows Update

By default, updated device drivers uploaded to Windows Update are downloaded and installed automatically on a client computer. You can amend this behavior through the Device Installation Settings dialog box, shown in Figure 4-7. The most straightforward method of accessing this dialog box is to type **device installation** in the Search text box on the Start menu and click Change Device Installation Settings.




FIGURE 4-7 Device Installation Settings dialog box

The default setting is Yes, Do This Automatically (Recommended). However, if you want to ensure that only device drivers you have tested are installed on the computers running Windows 7 in your organization, you can select Never Install Driver Software From Windows Update.

If you choose Always Install The Best Driver Software From Windows Update, Windows 7 will determine automatically whether a new driver is superior to one already installed. However, you will not be able to test the new driver before installation. Similarly, if you select Only Install

Driver Software From Windows Update If It Is Not Found On My Computer, you will not be able to test the new drivers before they are installed. In an enterprise environment, particularly if software is distributed through Windows Server Update Services (WSUS), driver updates from the Windows Update site are disabled and the site is removed from the search path.

Staging a Device Driver

If the device driver does not exist in the driver store, then an administrator needs to approve the device. This process is known as *staging*. You can configure a computer Group Policy so that ordinary users can approve the installation of a device in a specific device setup class, and you can stage a specific device driver so ordinary users can install that device. However, it would be most unwise to do this for all devices. You learn how to configure Group Policy to allow non-administrators to install specific devices and device setup classes in the practice session later in this lesson.

Windows 7 starts the staging process by searching for a matching driver package in folders specified by the DevicePath registry entry. You learn how to configure Windows 7 to search additional folders for device drivers in the practice session later in this lesson. If a suitable driver is not found, Windows 7 searches the Windows Update Web site. Finally, the user is prompted to insert installation media.

If a driver is found, the operating system checks that the user has permission to place the driver package in the driver store. The user must have administrator credentials or computer policy must be set to allow standard users to install the identified device. Windows 7 then checks that the driver package has a valid digital signature. If the driver package is unsigned or signed by a certificate not found in the Trusted Publishers store, Windows 7 prompts the user for confirmation. If the driver is approved by an authorized user, the operating system places a copy of the driver package in the driver store and installation continues.

Windows performs all the required security checks during staging, including the verification of administrator rights and the validation of digital signatures. After a driver package has been successfully staged, any user that logs on to that computer can install the drivers in the driver store by simply plugging in the appropriate device. There are no prompts, and no special permissions are required.

Quick Check

- You have four devices plugged into an unpowered USB hub. You have a powered hub with an empty slot and decide you could improve hardware performance by transferring one device from the first hub to the second. How do you determine which of the devices on the unpowered hub consumes the most power?

Quick Check Answer

- Double-click Universal Serial Bus Controllers, in Device Manager. Right-click the unpowered USB hub and choose Properties. On the Power tab, view the power required by each device in the Attached Devices list.

reference computer running Windows 7, you need to download the driver files from the vendor's Web site and then manually update the driver.

When you update a driver, the process is similar to the way a driver is installed when new hardware is added to the computer. If the driver is already authorized and in the driver store, or if an administrator has staged it, the update process, once started, proceeds automatically without user intervention. Otherwise, you can allow Windows 7 to search for drivers for the hardware device or you can specify a folder address manually as described in the previous section, "Installing a PnP Device." If a driver can be found that is more recent than that already installed, the administrator is prompted to approve the driver.

You can start the driver update process in Device Manager in two ways:

- Right-click the device and choose Update Driver Software.
- Double-click the device. On the Driver tab, click Update Driver.

You can install a new driver or one that has been updated. If, however, you consider the current driver installation might be corrupt, updating a driver will not by itself reinstall the current driver. In this case, you need to uninstall the driver and then go through the process of installing it again. As with updating a driver, you can either right-click the device in Device Manager and choose Uninstall, or you can double-click the device and click Uninstall on the Driver tab. The Driver tab is shown in Figure 4-11.




FIGURE 4-11 The Driver tab

Sometimes, rather than uninstalling and reinstalling a driver, you can solve problems by disabling it. If you have a driver conflict, you probably need to disable one of the drivers or stop it if it is a non-PnP device.

You cannot roll back a driver by right-clicking the device in Device Manager. You need to access the device's Driver tab. Unless a previous driver has been installed, the Roll Back

denies execute, read, and write permission to all types of removable media and overwrites any other configured removable media access rights policies. If you enable the All Removable Storage: Allow Direct Access policy in remote sessions, remote users can access removable storage devices in remote sessions.

Detailed step-by-step instructions to disable downloads to USB flash memory are given in the practice later in this lesson.

Changing Disk Type and Partition Style

Chapter 2, “Configuring System Images,” introduced the Disk Management tool and discussed the new features of the tools that let you create, attach, mount, detach, and delete a VHD file. In addition to these new features, you can use Disk Management to convert disks from basic to dynamic (and less typically, from basic to dynamic) and change the partition style. Windows 7 also provides the Diskpart command-line tool for disk management.

NOTE FDISK

Windows 7 does not support the Fdisk tool that was used for disk management in earlier versions of Windows.

You need Administrator or Backup Operator credentials to perform most Disk Management tasks. You can use the tool for managing hard disks and the volumes or partitions that they contain. You can initialize disks, create volumes, and format volumes with the FAT, FAT32, or NTFS file systems. Disk Management lets you perform disk-related tasks without needing to restart the system and most configuration changes take effect immediately.

Working with Partitions

Disk Management enables you to extend and shrink partitions directly from the interface. When you right-click a volume in the Disk Management console, you can choose whether to create a basic, spanned, or striped partition directly from the menu. If you add more than four partitions to a basic disk, you are prompted to convert the disk to dynamic or to the GPT partition style.

Disk Management allows you to change disks between various types and partition styles. However, some of the operations are irreversible (unless you reformat the drive). You should consider carefully the disk type and partition style that is most appropriate for your application. You can convert between the following partition styles:

- MBR can be converted to GPT if there are no volumes on the disk.
- MBR can be converted to dynamic, but the disk may become unbootable.
- GPT can be converted to MBR if there are no volumes on the disk.
- GPT can be converted to dynamic, but the disk may become unbootable.
- Dynamic can be converted to MBR if there are no volumes on the disk.
- Dynamic can be converted to GPT if there are no volumes on the disk.

3. On the Select Disks page, select from the available disks and then click Add to add the disks to the striped volume. Specify the amount of space to use on the disks for the striped volume.
4. On the Assign Drive Letter Or Path page, the default is to assign the next available drive letter to the new volume. You can also mount the volume on an empty NTFS folder on an existing volume.
5. On the Format Volume page of the New Striped Volume Wizard, choose the formatting options for the new volume. Windows 7 supports only NTFS formatting from the Disk Management snap-in. To format with FAT or FAT32, you need to use Diskpart.
6. On the Summary page, read the information provided. If this is satisfactory, click Finish.

To create a striped volume on a dynamic disk, enter a command with the following syntax at the DISKPART> prompt:

```
create volume stripe [size=<n>] disk=<n>[,n[,..]]
```

The total size of the stripe volume is the size multiplied by the number of disks.

Creating a Mirrored Volume (RAID-1)

A mirrored or RAID-1 volume provides availability and fault tolerance but does not improve performance. It uses two disks (or two portions on separate disks) that are the same size. Any changes made to the first disk of a mirror set are also made to its mirror disk. If the first disk fails, the mirror is broken and the second disk is used until the first is repaired or replaced. The mirror is then re-created, and the information on the working disk is mirrored on the repaired disk. The disadvantage of RAID-1 is that you need (for example) two 200-GB disks to hold 200 GB of data. The advantage is that you can mirror a system disk containing your operating system.

You create a mirrored volume using a very similar procedure to the one that creates a striped volume, except that you right-click the first disk of your mirror and click New Mirrored Volume to start the appropriate wizard. You then select the second disk. The second disk needs to have a portion of unallocated space that is at least as large as the disk you want to mirror. The drive letter for a mirrored volume is the same as the drive letter of the first disk.

You can also use the Diskpart tool to create a mirrored volume. At the DISKPART> prompt you first use the *select disk* command to select the first disk. You then enter a command with the syntax *add disk=<n>* to specify the mirror disk.

Creating a Striped Volume with Parity (RAID-5)

A striped volume with parity offers high availability, failover protection, and performance improvement. It requires at least three disks, or equally sized portions of unallocated space on at least three separate disks. The volume is striped in a similar way to RAID-0, but on each disk, some of the capacity is used to store parity information, which is compressed information about the contents of the other disks in the set.

- Windows 7 supports basic disks, dynamic disks, the MBR partition type, and the GPT partition type and allows you to convert from one to the other.
- Windows 7 offers software RAID-0 and RAID-1 volumes. You can also create simple and spanned volumes. You can shrink or expand a volume without needing to use third-party tools.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 2, “Managing Disks.” The questions are also available on the companion DVD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the “Answers” section at the end of the book.

1. Which Diskpart command converts an MBR disk to a GPT disk?
 - A. *convert gpt*
 - B. *convert mbr*
 - C. *convert basic*
 - D. *convert dynamic*
2. You require fault tolerance for your operating system so that your computer running Windows 7 Home Premium can still boot up if a disk fails. You have two disks and unallocated space on your second disk. What do you do?
 - A. Create a VHD and install an image of your computer on the VHD. Use BCDEdit to make the VHD bootable.
 - B. Create a RAID-0 volume.
 - C. Create a RAID-1 volume.
 - D. Create a RAID-5 volume.
3. You want to prohibit read, write, and execute access to all types of external storage devices. What computer policy setting do you enable?
 - A. All Removable Storage: Allow Direct Access In Remote Sessions
 - B. All Removable Storage Classes: Deny All Access
 - C. Removable Disks: Deny Read Access
 - D. Removable Disks: Deny Write Access
4. You are using the Diskpart tool to create a RAID-0 volume from unallocated space on Disks 1, 2, and 3. You want the volume to be as large as possible. What command do you enter?
 - A. `create volume stripe size=0 disk=1,2,3`
 - B. `create volume stripe disk=1,2,3`

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. These scenarios set up real-world situations involving the topics of this chapter and ask you to create a solution.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- If a device is not PnP, you need to supply administrator credentials to install it. You can prestage a device driver and (if necessary) digitally sign it so non-administrators can install it.
- You can prevent drivers downloading from Windows Update and installing automatically. You can also remove the Windows Update site from the search path for device drivers not in the device driver store. You can update, disable (or stop), uninstall, or roll back device drivers.
- Windows 7 enables you to manage disks, partitions, and volumes and to control access to removable devices. You can convert one disk type to another and one partition type to another. You can shrink or expand volumes.
- Windows 7 supports single, spanned, RAID-0 and RAID-1 volumes.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- **defragmentation**
- **driver store**
- **staging**
- **Redundant Array of Independent Disks (RAID)**
- **Trusted Publisher store**




FIGURE 5-12 Software Restriction Policy security levels

Enforcement

You can use the Enforcement Properties policy, shown in Figure 5-13, to specify whether Software Restriction Policies apply to all software files except libraries, such as DLLs, or all software files, including DLLs. If the default level is set to Disallowed and you configure the enforcement policies to apply to all software files, you need to configure rules for all the DLL files used by a program to use that program. Microsoft recommends that you do not include DLLs unless you are managing computers in a highly secure environment. This is primarily because managing rules for DLLs adds significantly to the amount of work that an administrator has to undertake to maintain Software Restriction Policies successfully.



FIGURE 5-13 Software Restriction Policy enforcement

You can use the Enforcement policy to apply Software Restriction Policies to all users or all users except for members of the local administrators group. You can also use this policy

Lesson 1: Configuring IPv4

As an IT professional with at least one year's experience, you will have come across IPv4 addresses, subnet masks, and default gateways. You know that in the enterprise environment, Dynamic Host Configuration Protocol (DHCP) servers configure IPv4 settings automatically and Domain Name System (DNS) servers resolve computer names to IPv4 addresses.

You might have configured a small test network with static IPv4 addresses, although even the smallest of modern networks tend to obtain configuration from a cable modem or a WAP, which in turn is configured by an Internet service provider (ISP). You might have set up Internet Connection Sharing in which client computers access the Internet through, and obtain their configuration from, another client computer.

You have probably come across Automatic Private Internet Protocol (APIPA) addresses that start with 169.254 when debugging connectivity because computers that fail to get their IPv4 configuration addresses from DHCP typically configure themselves using APIPA instead—so an APIPA address can be a symptom of DHCP failure or loss of connectivity, although it is also a valid way of configuring isolated networks that do not communicate with any other network, including the Internet.

However, you might not have been involved in network design or have subnetted a network. Subnetting is not as common these days, when private networks and NAT give you a large number of addresses you can use. It was much more common in the days when all addresses were public and administrators had to use very limited allocations. Nevertheless, subnetting remains a useful skill and subnet masks are likely to be tested in the 70-680 examination.

In this lesson, you look at the tools available for manipulating IPv4 addresses and subnet masks and implementing IPv4 network connectivity. The lesson considers the Network And Sharing Center, the Netstat and Netsh command-line tools, Windows Network Diagnostics, how you connect a computer to a network, how you configure name resolution, the function of APIPA, how you set up a connection for a network, how you set up network locations, and how you resolve connectivity issues.

Before you look at all the tools for manipulating and configuring IPv4, you first need to understand what the addresses and subnet masks mean. You will learn the significance of addresses such as 10.0.0.21, 207.46.197.32, and 169.254.22.10. You will learn why 255.255.255.128, 255.255.255.0, 255.255.254.0, and 255.255.252.0 are valid subnet masks, whereas 255.255.253.0 is not. You will learn what effect changing the value of the subnet mask has on the potential size of your network and why APIPA addresses do not have default gateways.

This chapter starts with an introduction to IPv4, in particular IPv4 addresses, subnet masks, and default gateways. It continues with the practical aspects of configuring and managing a network.

So what identifies the computer and what identifies the subnet? To discover this, we need to look at the next value, the subnet mask. Subnet masks are most peculiar numbers. They represent binary numbers that consist of all ones followed by all zeros. For example:

255.255.255.0 is the binary number 11111111 11111111 11111111 00000000.

The actual value of this number is irrelevant. What matters is the number of ones and zeros. A one says that the corresponding bit in the IPv4 address is a network address bit. A zero says that the corresponding bit in the IPv4 address is a computer or host address bit.

In the example given, the last 8 bits of the subnet mask are all zero. So the host address is the final octet of the subnet address, or 143. The network address of the subnet is 10.16.10.0. Because hosts are defined by a single octet in this example, the 10.16.10.0 subnet contains 254 host addresses. The first IPv4 address in the subnet is 10.16.10.1. The last is 10.16.10.254. The number 10.16.10.0 identifies the subnet and is called the *subnet address*. The number 10.16.10.255 is called the *broadcast address* and is used when a packet needs to be sent to every host on a subnet.

Subnetting and Supernetting

You can split a subnet into smaller subnets by adding ones to the end of the ones in the subnet mask. If you have two (or more) suitable contiguous subnets, you can merge them into a single subnet by changing one or more ones at the end of the ones in the subnet masks to zeros. These techniques are known as *subnetting* and *supernetting*, respectively.

If an organization has a significant number of computers on its network (say over 100—this number varies depending on the type, volume, and pattern of traffic on the network) or if it has several geographic locations, the organization probably creates several subnets. If a subnet contains too many computers and other devices, it tends to slow down because there is a greater chance of two computers trying to put data onto the network simultaneously and causing a collision. Dividing a network into several subnets reduces the likelihood of such collisions.

At the router that connects to the Internet, however, the organization uses supernetting to combine (or summarize) the subnets so that they can be defined with a single network address that will be translated to a public address on the Internet. Public addresses and address translation are discussed later in this lesson.

MORE INFO SUBNETTING AND SUPERNETTING

For more information about supernetting and subnetting, and about CIDR and VLSM technologies, see <http://support.microsoft.com/kb/164015>.

NOTE CIDR NOTATION

Because the subnet mask 255.255.255.0 consists of 24 ones followed by 8 zeros, you can also write it as /24. A subnet with a network address 192.168.0.0 and a subnet mask 255.255.255.0 (for example) is then designated 192.168.0.0/24. This is sometimes called *CIDR notation*. A subnet mask with 25 ones followed by 7 zeros is a /25 subnet mask. In dotted decimal, this would be 255.255.255.128.

The final value shown in Figure 6-1 is the *default gateway*. This is the IPv4 address of the router connection on the same subnet as the IPv4 address of the host computer. If an IPv4 packet has a destination address of a different subnet, it is routed through other subnets via the router until it finds the destination it is looking for. If you browse to a Web site, for example, you need to send data to the Web server for that site, which has an IPv4 address somewhere on the Internet.

Put simply, some packets need to get out of your subnet and go to another network (for example, the Internet). Your computer sends these packets to a routing device. This can be a hardware router, a server that is configured as a router, or the computer or wireless router through which the other computers in a small office/home office (SOHO) network access the Internet. The default gateway is the address within the subnet of the routing device (which has at least one more IPv4 address on another subnet). It is where outgoing packets leave the subnet. It is also where incoming packets from other networks enter the subnet.

✓ Quick Check

1. What is the binary number 00001010 11110000 10101010 01000000 in dotted decimal notation?
2. Are the IPv4 addresses 192.168.1.200 and 192.168.1.24 on the same subnet? Both have a subnet mask of 255.255.255.0.
3. Is 10.0.0.130 a valid IPv4 address on the 10.0.0.0/25 subnet?

Quick Check Answers

1. 10.240.170.64.
2. The subnet mask specifies that the final octet holds the host address. Therefore the first three octets hold the subnet's network address. In both cases, this is 192.168.1.0, so the computers are on the same subnet.
3. No. The /25 subnet mask specifies 25 ones and therefore $32 - 25 = 7$ zeros. Zeros denote host address. Therefore, the host address is from 0000001 to 1111110 binary (0000000 is the network address and 1111111 is the broadcast address). In decimal, this is 1 to 126. So the valid IPv4 addresses on the network are 10.0.0.1 to 10.0.0.126. 10.0.0.130 is not in this range and therefore is not valid on this subnet. It is an address on another subnet (for example, 10.0.0.128/25).

Network Services

IPv4 configuration and operation relies on a number of network services. In an enterprise environment, these services (apart from APIPA) are implemented on servers. However, on a small network, DHCP and DNS services can be provided by a client running ICS or by a WAP. Services associated with IPv4 include the following:

(URL) into a browser—the information that the request returns (the Web page) needs to find its way back to your computer, which has an internal IPv4 address on your local area network (LAN). Typically, your ISP allocates only one public IPv4 address that all the computers on your LAN share when accessing the Internet. NAT deals with this situation and ensures that IPv4 packets from the Internet reach the correct LAN destinations.

MORE INFO NETWORK ADDRESS TRANSLATION

For more information about NAT, see <http://technet.microsoft.com/en-us/library/cc739385.aspx>.

Public and Private IPv4 Addresses

Every device on the Internet has its own unique public IPv4 address that is shared with no other device (a LAN also has at least one IPv4 address that is unique on the Internet). For example, if you type a URL such as **http://www.adatum.com** into your Web browser, the FQDN *www.adatum.com* identifies a Web server that has a public IPv4 address—for example, 207.46.197.32.

Any organization that has an Internet presence is allocated one or more public IPv4 addresses that that organization and only that organization can use. The Internet Assigned Numbers Authority (IANA) issues and controls public IPv4 addresses through various agencies—for example, the United Kingdom Education and Research Network (UKERNA). In the case of a SOHO network, the ISP allocates one public IPv4 address from a range that IANA or one of its agencies has allocated to the ISP.

Most organizations do not have enough public IPv4 addresses to allocate one to every device on their networks. Also, issuing public IPv4 addresses to computers in an organization's network has security implications. Instead, organizations use private IPv4 addresses for their internal networks and use NAT to translate these addresses to a public address or addresses for Internet access.

Private IPv4 addresses should never be used on the Internet, and typically a router on the Internet ignores private IPv4 addresses. An organization can use whatever private IPv4 address range it chooses without requiring permission from IANA. Because private IPv4 addresses are internal to an organization, many organizations can use the same range of IPv4 addresses without causing IPv4 conflicts. Most computers on internal networks do not need a unique public address but instead share a single public address that identifies their LAN and that NAT translates to their private addresses. Only devices on a LAN that have an Internet presence—for example, Web servers, e-mail servers, and DNS servers—require a unique public address mapped through NAT to their internal private address.

IANA has reserved the following three blocks of IPv4 address space for private networks:

- 10.0.0.0/8 (10.0.0.1 through 10.255.255.254)
- 172.16.0.0/12 (172.16.0.1 through 172.31.255.254)
- 192.168.0.0/16 (192.168.0.1 through 192.168.255.254)

The syntax of the *Netstat* command is as follows:

```
netstat [-a] [-e] [-n] [-o] [-p Protocol] [-r] [-s] [Interval]
```

The parameters implement the following functions:

- **-a** Displays all active connections and the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports on which the computer is listening.
- **-e** Displays Ethernet statistics, such as the number of bytes and packets sent and received. This parameter can be combined with *-s*.
- **-n** Displays active connections. Addresses and port numbers are expressed numerically and no attempt is made to determine names.
- **-o** Displays active connections and includes the process ID (PID) for each connection. You can find the application based on the PID on the Services tab in Windows Task Manager. This parameter can be combined with *-a*, *-n*, and *-p*.
- **-p protocol** Shows connections for the protocol specified by the protocol variable, which can be *tcp*, *udp*, *tcpv6*, or *udpv6*. If this parameter is used with *-s* to display statistics by protocol, which can be *tcp*, *udp*, *icmp*, *ip*, *tcpv6*, *udpv6*, *icmpv6*, or *ipv6*.
- **-s** Displays statistics by protocol. By default, statistics are shown for the TCP, UDP, ICMPv4, ICMPv6, IPv4, and IPv6 protocols. The *-p* parameter can be used to specify a set of protocols.
- **-r** Displays the contents of the IP routing table. This is equivalent to the *route print* command.
- **interval** Displays the selected information periodically. The number of seconds between each display is defined by the interval parameter. If this parameter is omitted, *Netstat* prints the selected information only once.

Netstat provides statistics for the following:

- The name of the protocol (TCP or UDP)
- The IP address of the local computer and the port number being used
- The IP address and port number of the remote computer
- The state of a TCP connection

MORE INFO TCP CONNECTION STATES

For more information about the states of a TCP connection, see <http://support.microsoft.com/kb/137984>. This article was written some time ago but remains relevant to Windows 7.

For example, to display both the Ethernet statistics and the statistics for all protocols, enter the following command:

```
netstat -e -s
```

To display the TCP statistics for the IPv4 protocol, enter the following command:

```
netstat -s -p tcp
```

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 1, “Configuring IPv4.” The questions are also available on the companion DVD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the “Answers” section at the end of the book.

1. Which command-line command displays the IP configuration of a computer’s interfaces?
 - A. *Ping*
 - B. *Tracert*
 - C. *Ipconfig*
 - D. *Netstat*
2. Which of the following methods can you use to display the properties of a LAN connection? (Choose all that apply.)
 - A. In the Network And Sharing Center, click Internet Options. On the Connections tab of the Internet Properties dialog box, click LAN Settings.
 - B. In the Network And Sharing Center, click Change Adapter Settings. Right-click the LAN connection and choose Status. In the Local Area Connection Status dialog box, click Properties.
 - C. In the Network And Sharing Center, click Change Adapter Settings. Right-click the LAN connection and choose Properties.
 - D. In the Network And Sharing Center, click Change Adapter Settings. Double-click the LAN connection. In the Local Area Connection Status dialog box, click Properties.
3. You are configuring static IPv4 addresses for two computers, Perth and Brisbane, on an isolated private wired subnet. You configure Perth with the IPv4 address 172.16.10.140 and the subnet mask 255.255.255.0. You configure Brisbane with the IPv4 address 172.16.10.210 and the subnet mask 255.255.255.0. You enter **ping 172.16.10.140** on Brisbane, but the command times out. Similarly, entering **ping 172.16.10.210** on Perth fails to locate the Brisbane computer’s IPv4 address. What is the likely reason for this lack of connectivity?
 - A. DNS service is not available on the subnet.
 - B. The computers should have different subnet masks.
 - C. You have not specified a default gateway.
 - D. You need to permit ICMPv4 traffic through the firewalls of both computers.

4. You have created the statically configured wired subnet 10.0.10.128/25. Currently the only device on the subnet is a router with the IPv4 address 10.0.10.129. You plug a computer into an Ethernet port on the subnet. What command configures the computer correctly on the subnet?
- A. `netsh interface ipv4 set address "local area connection" static 10.0.10.162 255.255.255.0 10.0.10.129`
 - B. `netsh interface ipv4 set address "local area connection" static 10.0.10.162 255.255.255.128 10.0.10.129`
 - C. `netsh interface ipv4 set address name="local area connection" dhcp`
 - D. `netsh interface ipv4 set address name="local area connection" static 10.0.10.16 255.255.255.128 10.0.10.129`
5. You want to examine the contents of both the IPv4 and the IPv6 route table. What command do you use? (Choose all that apply.)
- A. `netsh interface ipv4 show route`
 - B. `tracert -d`
 - C. `route print`
 - D. `netstat -r`
 - E. `netstat -a`

MORE INFO Netsh

Netsh is an exceptionally powerful and versatile utility that enables you to carry out a very large number of configuration tasks through a command-line interface. For more information, see <http://technet.microsoft.com/en-us/library/cc785383.aspx>.

Quick Check

- What *Netsh* command lists site IDs?

Quick Check Answer

- `netsh interface ipv6 show address level=verbose`

Verifying IPv6 Connectivity

To verify connectivity on a local network, your first step should be to flush the neighbor cache, which stores recently resolved link-layer addresses and might give a false result if you are checking changes that involve address resolution. You can check the contents of the neighbor cache by entering **netsh interface ipv6 show neighbors**. Entering **netsh interface ipv6 delete neighbors** flushes the cache. You need to run the command prompt as an administrator to use these commands.

You can test connectivity to a local host on your subnet and to your default gateway by using the *Ping* command. Note that Windows Firewall blocks *Ping* commands by default and you need to allow ICMPv6 packets through the firewalls of both computers before one can ping the other by its IPv4 address. You can add the interface ID to the IPv6 interface address to ensure that the address is configured on the correct interface. Figure 6-22 shows a *Ping* command using an IPv6 address and an interface ID.

To check connectivity to a host on a remote network, your first task should be to check and clear the destination cache, which stores next-hop IPv6 addresses for destinations. You can display the current contents of the destination cache by entering **netsh interface ipv6 show destinationcache**. To flush the destination cache, **enter netsh interface ipv6 delete destinationcache**. As before, these commands need administrator credentials.

Your next step is to check connectivity to the default router interface on your local subnet. This is your default gateway. You can identify the IPv6 address of your default router interface by using the *Ipconfig*, *Netsh interface ipv6 show route*, or *Route print* command. You can also specify the zone ID, which is the interface ID for the default gateway on the interface on which you want the ICMPv6 Echo Request messages to be sent. When you have ensured that you can reach the default gateway on your local subnet, ping the remote host by its IPv6 address. Note that you cannot ping a remote host (or a router interface) by its link-local IPv6 address because link-local addresses are not routable.

If you can connect to the default gateway but cannot reach the remote destination address, trace the route to the remote destination by using the *Tracert -d* command followed by the destination IPv6 address. The *-d* command-line switch prevents the Tracert tool from performing a DNS reverse query on router interfaces in the routing path. This speeds up the display of the routing path. If you want more information about the routers in the path, and particularly if you want to verify router reliability, use the *Pathping* command, again followed by the destination IPv6 address.

✓ Quick Check

- What *Netsh* command could you use to identify the IPv6 address of your default router interface?

Quick Check Answer

- `netsh interface ipv6 show route`

Troubleshooting Connectivity

If you cannot connect to a remote host, you first need to check the various hardware connections (wired and wireless) in your organization and ensure that all network devices are up and running. If these basic checks do not find the problem, the Internet Protocol Security (IPSec) configuration might not be properly configured, or firewall problems (such as incorrectly configured packet filters) might exist.

You can use the IP Security Policies Management console, shown in Figure 6-26, to check and configure IPSec policies and the Windows Firewall With Advanced Security console (shown previously in Figures 6-11 and 6-12 in Lesson 1) to check and configure IPv6-based packet filters.




FIGURE 6-26 The IP Security Policies Management console

Wireless Network Technologies

Advantages of wireless networks include mobility and easy physical installation (you do not need to run cables under the floor). Disadvantages include a slower connection (typically) than a wired network and interference from other wireless devices, such as cordless phones.

Currently there are (arguably) four types of wireless network technologies in common use:

- **802.11b** Up to 11 megabits per second (Mbps); good signal range; low cost. This technology allows fewer simultaneous users than the other options and uses the 2.4-gigahertz (GHz) frequency. This frequency is prone to interference from microwave ovens, cordless phones, and other appliances.
- **802.11a** Up to 54 Mbps; more simultaneous users than 802.11b, but a smaller signal range; expensive. This option provides a fast transmission speed and uses the 5-GHz frequency, which limits interference from other devices. However, its signal is more easily obstructed by walls and other obstacles and it is not compatible with 802.11b network adapters, routers, and access points.
- **802.11g** Up to 54 Mbps (under optimal conditions); more simultaneous users than 802.11b; very good signal range; not easily obstructed. This option is compatible with 802.11b network adapters, routers, and access points, but it uses the 2.4-GHz frequency and has the same interference problems as 802.11b. It is also more expensive than 802.11b.
- **802.11n** Still in draft format, although this situation may have changed by the time you read this book. However, a number of vendors are manufacturing equipment using the current draft 802.11n standard. Most 802.11n devices are compatible with 802.11b and 802.11g. 802.11n builds on previous 802.11 standards by adding multiple-input, multiple-output (MIMO), which uses multiple transmitter and receiver antennas to improve the system performance.

802.11b is adequate for most home and many small-office applications. If, however, your network carries a high volume of streaming media (video or music) traffic, or if interference is a major problem, you might consider 802.11a. If you already have 802.11b devices on your network but require high-speed transmission between specified network points, you might consider 802.11g. Most modern WAPs available from computer equipment retailers now are 802.11g.

If you have more than one wireless network adapter in your computer, or if your adapter uses more than one standard, you can specify which adapter or standard to use for each network connection.



EXAM TIP

Several 802.11 standards exist in addition to 802.11a, 802.11b, and 802.11g. However, the standards described in this lesson are those in common use. If you see any other standard (for example, 802.11d) given as a possible answer in the examination, that answer is almost certainly wrong.

Lesson 1: Managing Windows Firewall

The firewall that ships with Windows 7 is designed to keep your computer safe. It will help keep your computer safe when it is connected to the protected network in the office or the less-safe public WiFi networks of coffee shops and airport lounges. In this lesson, you learn the differences between Windows Firewall and Windows Firewall with Advanced Security. You also learn about *connection security rules*, which you can use to limit your computer's network communication so that it occurs only with other computers that have proven their identity.

After this lesson, you will be able to:

- Configure rules for multiple profiles.
- Allow or deny applications.
- Create network-profile-specific rules.
- Configure notifications.
- Configure authenticated exceptions.

Estimated lesson time: 40 minutes

Windows Firewall

Firewalls restrict network traffic based on a collection of configurable rules. Another name for these rules is *exceptions*. When traffic reaches a network interface protected by a firewall, the firewall analyzes it, either discarding the traffic or allowing it to pass on the basis of the rules that have been applied to the firewall. Windows 7 uses two firewalls that work together: Windows Firewall and the Windows Firewall with Advanced Security (WFAS). The primary difference between these firewalls relates to the complexity of the rules that can be configured for them. Windows Firewall uses simple rules that directly relate to a program or service. WFAS allows for more complicated rules that filter traffic on the basis of port, protocol, address, and authentication. WFAS will be covered in more detail later in this lesson.

When thinking about how firewall rules work, remember that unless a rule exists that explicitly allows a particular form of traffic, the firewall will drop that traffic. In general, you must explicitly allow traffic to pass across a firewall, though there will be some occasions when you need to configure a deny rule. You will learn about deny rules later in this lesson. Windows Firewall and WFAS ship a minimum number of default rules that allow you to interact with networks. This means that although you are able to browse the Web without having to configure a firewall rule, if you try to use an application to interact with the network that is not covered by a default rule, such as File Transfer Protocol (FTP), you receive a warning. This behavior is different to earlier versions of Microsoft Windows, such as Windows XP, where the firewall blocked only incoming traffic and did not block outgoing traffic. The firewall in Windows 7 blocks most outbound traffic by default. When a program is blocked for the first time, you are notified by the firewall, as shown in Figure 7-1, allowing you to configure an exception that allows traffic of this type in the future.

extension. Exported policies use a binary format, not Extensible Markup Language (XML) format like many other Windows 7 configuration files. You can also export and import firewall policies in the same .wfw format using the *netsh advfirewall export* or *netsh advfirewall import* commands.

Managing WFAS with Netsh

You can use the Netsh.exe command-line utility from an elevated command prompt to manage WFAS rules. The advantage of this is that you can combine it with Windows Remote Shell (WinRS), which you will learn about in the next lesson, to manage WFAS rules on other computers running Windows 7 on your network. You can also use Netsh.exe to script the creation of firewall rules on stand-alone computers that are not members of an AD DS domain and hence are not subject to domain-applied Group Policy.

To use Netsh.exe to create WFAS firewall rules, you need to be in the *advfirewall firewall* context. The following are some examples of using WFAS to create firewall rules:

- To create a rule named WebServerRule that applies in the domain profile and allows inbound traffic on TCP port 80, issue the command **netsh advfirewall firewall add rule name="WebServerRule" profile=domain protocol=TCP dir=in localport=80 action=allow.**
- To create a rule named AllowCalc that allows inbound traffic to the Calc.exe application in all network profiles, issue the command **netsh advfirewall firewall add rule name="Calc" dir=in program="c:\windows\system32\calc.exe" action=allow.**
- To create a rule named BlockFTP that blocks outbound traffic from the Ftp.exe application, issue the command **netsh advfirewall firewall add rule name="BlockFTP" dir=out program="c:\windows\system32\ftp.exe" action=block.**



EXAM TIP

Know when you need to use WFAS to create a rule and when you can use Windows Firewall.

PRACTICE Configuring Windows Firewall

In this practice, you use Windows Firewall and WFAS to configure two different firewall rules. By interacting with the interface, you learn more about the types of rules that you can configure with each tool.

EXERCISE 1 Configuring Firewall Rules with Windows Firewall

In this exercise, you configure a rule that allows incoming traffic to the Internet Backgammon application. To complete this exercise, perform the following steps:

1. Log on to Canberra with the Kim_Akers user account.
2. Click Start, Control Panel, and System And Security.

You can manage all shared folders on a client running Windows 7 centrally using the Shared Folders node of the Computer Management console. The Shares node, shown in Figure 8-11, displays all shared folders on the computer. The Sessions node provides details on which remote users currently are connected to shared folders, where they are connecting from and how long they have been connected. The Open Files node displays the folders and files that remote users are accessing. You can edit the properties of an existing share by right-clicking it within this console and selecting properties. You can create a shared folder by right-clicking the Shares node and then clicking New Share. This starts the Create A Shared Folder Wizard. You use this wizard to create a shared folder in a practice exercise at the end of this lesson.




FIGURE 8-11 Viewing shares

The *Net Share* command allows for management of shared folders from the command line. You can script this command to automate the creation of shared folders on clients running Windows 7. To create a shared folder, use the command:

```
net share sharename=drive:path
```

To assign permissions when creating the shared folder, use the command:

```
net share sharename=[path] /grant:user,[Read/Change/Fu11]
```

(for example `net share tempus=c:\temp /grant:Bob,Change`)

You can also use the *Net Share* command to configure caching options as well as limit the number of users that can connect to the shared folder. You can view the properties of a shared folder by running the command

```
net share sharename
```

as shown in Figure 8-12. You can view the properties of all shared folders, including which directories are associated with particular folders, by using the *Net Share* command without any options.

When you move a file from a folder on one volume to a folder on another volume, the file behaves the same way that it does when you copy it and inherits the permissions of the destination folder. The same applies to a folder. If you move a folder from one volume to another, that folder and all its contents inherit the permissions assigned to the destination folder.

Robocopy.exe is a command-line utility that is included with Windows 7 that allows you to copy files while retaining their existing NTFS permissions. You can also use Robocopy.exe to move files from one volume to another while allowing them to retain their permissions. You should consider Robocopy.exe to be an exception to the normal rules of copying and moving files. In an exam situation, you should assume that the normal rules apply unless the question mentions Robocopy.exe. To use Robocopy.exe to copy all files and folders from the folder name C:\Example\ to the folder D:\Destination, use the command

```
Robocopy.exe c:\example d:\destination /copyall /e
```

NOTE MOVING TO FAT VOLUMES

If you move a file or folder to a volume formatted with the FAT or FAT32 file system, all NTFS permissions are lost.

Combined Share and NTFS Permissions

When a user accesses a file hosted on a shared folder, both the share permissions, which you learned about in Lesson 1, and the NTFS permissions apply. The most restrictive permission of the share and the NTFS permissions apply. For example, if a group is assigned the Read permission at the Share level and the Modify permission through file and folder permissions, the user has only Read access to files and folders when connecting to the shared folder over the network. Similarly, if a user has Full Control access at the share level and Read access assigned to the folder through NTFS permissions, the user has only Read access and is unable to modify or delete files and folders hosted on the share.

Configuring Auditing

Auditing allows you to monitor which users and groups access specific files and folders. You most likely do not want to monitor who accesses every document in your organization; you are most likely to use auditing only on sensitive documents. For example, you would use auditing to track who accessed the spreadsheet containing employee salaries, but you would not use auditing to track who accessed the break room cleanup roster. Auditing can tell you who opened a document, who modified a document, and who tried to open a document and failed. You can audit the use of any of the special permissions listed in Table 8-2. You can perform auditing only on volumes that are formatted using the NTFS file system.

The audit policies in Windows 7 allow a greater degree of granularity in tracking audit events compared to the audit policies in previous versions of Windows. For example, in Windows XP, you could audit nine broad event categories: in Windows 7, there are 53 different event categories. This allows you to be more specific about the types of events you

audit. To configure auditing to track which users access specific files and folders on clients running Windows 7, do the following:

1. Open the Local Group Policy Editor and navigate to the Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options node and set the Audit: Force Audit Policy Subcategory Settings (Windows Vista Or Later) To Override Audit Policy Category Settings policy to Enabled.
2. In the Local Group Policy Editor, navigate to the Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy Configuration\System Audit Policies – Local Group Policy Object\Object Access Node and set the Audit File System policy, as shown in Figure 8-27.




FIGURE 8-27 Configuring audit policies

3. Edit the properties of the file or folder that you wish to audit. On the Security tab, click Advanced, then click the Auditing tab, and then click Continue to elevate privileges.
4. Click Add and add the groups for which you want to audit access. If you want to audit the access of all users, select the Everyone group. Once you have selected the security group, you must select which of the special privileges you want to Audit. Figure 8-28 shows an auditing configuration to track successful file reads, writes, and deletes.
5. Auditing events will now be written to the Security log, which can be accessed using Event Viewer.

14. Log off as Jeff_Phillips and resume the Kim_Akers session. The User Access To Encrypt.txt dialog box should still be present on the screen because you switched to the other account and left the existing session active in memory.
15. In the User Access To Encrypt.txt dialog box, click Add. Verify that there are two encryption certificates present in the Windows Security dialog box. Click the Jeff_Phillips certificate, as shown in Figure 8-31, and then click OK.



FIGURE 8-31 Additional EFS certificate available

16. Click OK three times to close the Properties dialog box.

EXERCISE 2 Exploring File and Folder Permissions

In this exercise, you explore how file and folder permissions vary when you copy and move files between two folders. You use the Icacls and Effective Permissions tools during this exercise.

1. If you have not done so already, log on to Canberra using the Kim_Akers user account.
2. Open an elevated command prompt and issue the following commands:

```
net localgroup Research /add
net localgroup Accounting /add
net localgroup Research Jeff_Phillips /Add
net localgroup Accounting Jeff_Phillips /Add
mkdir c:\source
mkdir c:\destination
icacls c:\source /grant Research:(OI)(CI)M
icacls c:\destination /grant Accounting:(OI)(CI)RX
icacls c:\destination /deny Jeff_Phillips:(OI)(CI)W
```

Cache modes determine how the branch office cache functions. BranchCache can operate in one of two modes: Hosted Cache mode or Distributed Cache mode. You will learn about these modes during the rest of this lesson.

Hosted Cache Mode

Hosted Cache mode uses a centralized local cache that is hosted on a branch office server running Windows Server 2008 R2. You can enable the hosted cache server functionality on a server running Windows Server 2008 R2 that you use for other functions without a significant impact on performance. This is because if you found that files hosted at another location across the WAN were being accessed so frequently that there was a performance impact, you would use a solution like Distributed File System (DFS) to replicate them to the branch office instead of using BranchCache. The advantage of Hosted Cache mode over Distributed Cache mode is that the cache is centralized and always available. Parts of the distributed cache become unavailable when the clients hosting them shut down. You will learn more about Distributed Cache mode later in this lesson.

Hosted Cache mode requires a computer running Windows Server 2008 R2 be present and configured properly in each branch office. You must configure each BranchCache client with the address of the BranchCache host server running Windows Server 2008 R2.

When setting up the Hosted Cache mode server, it is necessary to do the following:

- Install the BranchCache feature.
- Install an Secure Sockets Layer (SSL) certificate where the subject name is set to the fully qualified domain name (FQDN) of the hosted cache server. This involves importing the SSL certificate into the Local Computer's certificate store, making note of the certificate thumbprint, and then binding the certificate using the command `netsh http add sslcert ipport=0.0.0.0:443 certhash=<thumbprint> APPID={d673f5ee-a714-454d-8de2-492e4c1bd8f8}`
- Ensure that all clients trust the certificate authority that issued the SSL certificate installed on the hosted cache server.

Hosted Cache mode is not appropriate for organizations that do not have their own Active Directory Certificate Services infrastructure or do not have the resources to deploy a dedicated server running Windows Server 2008 R2 to each branch office.

MORE INFO CONFIGURING HOSTED CACHE SERVERS

To learn more about configuring a Windows Server 2008 R2 server as a hosted cache server, including how to change the default ports used, consult the following document on TechNet: [http://technet.microsoft.com/en-us/library/dd637793\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd637793(WS.10).aspx).

EXERCISE 2 Configuring and Exporting UAC Policies

In this exercise, you configure User Account Control policies using the Local Security Policy editor.

1. If you have not done so already, log on to computer Canberra using the Kim_Akers user account.
2. Using Windows Explorer, create the C:\Export folder.
3. In the Search Programs And Files text box, type **Edit Group Policy**. Click the Edit Group Policy item.
4. Ensure that the Computer Configuration\Windows Settings\Security Settings node is selected. Open the Action menu and then choose Export Policy.
5. Save the exported policy as C:\Export\Base_policy.inf
6. Within Security Settings, select the Local Policies\Security Options node. Double-click the User Account Control: Behavior Of The Elevation Prompt For Administrators In Admin Approval Mode policy.
7. Select the Prompt For Credentials On The Secure Desktop setting, as shown in Figure 9-9, and then click OK.




FIGURE 9-9 Prompt For Credentials On The Secure Desktop

8. Click Start. In the In the Search Programs And Files text box, type **gpupdate /force** and press Enter.
9. Click Start. In the In the Search Programs And Files text box, type **User Accounts**. Click the User Accounts item on the Start menu.

in the Computer Configuration\Administrative Templates\Network\TCPIP Settings\IPv6 Transition Technologies node. You can see this collection of policies in Figure 10-1.




FIGURE 10-1 DirectAccess group policies

As computers must be domain-joined and be specifically added to security groups configured to support DirectAccess, there is little scope for local administrators to have to modify local policies on individual client running Windows 7 computers. When you configure DirectAccess on the DirectAccess server, it creates a GPO at the domain level and filters it for a specific security group. This GPO applies the following policies:

- **6to4 Relay Name** This policy sets the 6to4 relay name and is configured to use one of the public IPv4 addresses applied to the DirectAccess server.
- **IP-HTTPS State** This policy sets the Uniform Resource Locator (URL) of the IP-HTTPS server, which will be the FQDN of one of the public IPv4 addresses applied to the DirectAccess server. The default policy state uses IP-HTTPS as a connection of last resort. It is possible to set this policy to always use IP-HTTPS even if other connectivity options, such as 6to4 or Teredo, are available.
- **Teredo Default Qualified** This policy determines whether Teredo will be used. It is set to enabled for DirectAccess clients.
- **Teredo Server Name** This policy sets the address of the Teredo server. This address will be one of the public IPv4 address assigned to the DirectAccess server.

The final policy that is configured when you set up the DirectAccess server is the name resolution policy. The DirectAccess name resolution policy is located in the Computer Configuration\Policies\Windows Settings\Name Resolution Policy and is shown in Figure 10-2.

Case Scenario 1: Wingtip Toys DirectAccess

Wingtip Toys currently has 40 laptop computers running Windows Vista Business. Wingtip Toys wants to deploy DirectAccess because many of the users of these computers would prefer an automatic connection to the company network when they are in remote locations, rather than relying on a manual VPN connection. Wingtip Toys wants to replace their existing server running Windows Server 2003 R2 x64 Routing and Remote Access with a DirectAccess server. This server has two network cards and is assigned two consecutive public IPv4 addresses on the Internet interface. This server is a member of the WingtipToys.internal domain. The server has already been assigned the appropriate computer certificates.

With these facts in mind, answer the following questions:

1. What steps should Wingtip Toys take to create the DirectAccess server?
2. What type of group should you create to support DirectAccess?
3. What steps should you take to prepare client computers to use DirectAccess?

Case Scenario 2: Remote Access at Tailspin Toys

Tailspin Toys is deploying Windows 7 Enterprise to 300 laptop computers. You want to ensure that future VPN users will be able to stay connected to their VPN sessions if they switch from using a public Wi-Fi connection to using the cellular modem cards provided to them by the company. Users should be able to authenticate with their user names and passwords. Your existing VPN infrastructure uses NAP. The current Routing and Remote Access server is running the Windows Server 2008 x64 operating system. This system blocks VPN access to clients running Windows Vista Enterprise that do not have the most recent software updates or antivirus definitions installed. Presently, NAP blocks noncompliant clients from accessing the network. These clients cannot access the VPN until they connect to the corporate network directly and are able to download antivirus and software updates. You want to upgrade your quarantine network so that noncompliant clients can undergo remediation while connected remotely. Tailspin Toys has an Active Directory Certificate Services deployment.

With these facts in mind, answer the following questions:

1. What steps do you need to take to support VPN Reconnect at Tailspin Toys?
2. What additions should you make to the quarantine network so that clients can become compliant?
3. Which authentication protocol should you use for Tailspin Toys?

Suggested Practices

To help you master the exam objectives presented in this chapter, complete the following tasks.

Configuring a BitLocker DRA

Data Recovery Agents (DRAs) are special user accounts that you can use to recover encrypted data. You can configure a DRA to recover BitLocker-protected drives if the recovery password or keys are lost. The advantage of a DRA is that you can use it organization-wide, meaning that you can recover all BitLocker-encrypted volumes using a single account rather than having to recover a specific volume's recovery password or key.

The first step you must take in configuring BitLocker to support DRAs is to add the account of a DRA to the Computer Configuration\Windows Settings\Security Settings\Public Key Policies\BitLocker Drive Encryption node, as shown in Figure 11-3. A DRA account is a user account enrolled with a special type of digital certificate. In organizational environments, this digital certificate is almost always issued by an AD DS certificate authority (CA).




FIGURE 11-3 Assigning the recovery key

After you have configured the DRA, it is also necessary to configure the Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Prove The Unique Identifiers For Your Organization policy to support DRA recovery. BitLocker works with DRAs only when an identification field is present on a volume and the value matches that configured for the computer. Figure 11-4 shows this policy configured with the identification field set to ContosoBitLockerSelfHost. You also use this policy when denying write access to removable devices not protected by BitLocker. You will learn more about denying write access to removable devices later in this lesson.

After you have configured the DRA and the Identifiers, you need to configure the following policies to allow specific volume types to utilize the DRA:

- Choose How BitLocker-Protected Operating System Drives Can Be Recovered
- Choose How BitLocker-Protected Fixed Drives Can Be Recovered
- Choose How BitLocker-Protected Removable Drives Can Be Recovered

Advanced Power Plan Settings

To configure the advanced power plan settings, click the Change Advanced Power Settings item in the Edit Plan Settings dialog box. Unlike the basic plan settings, which a user who is not a member of the local administrators group can modify, only a user with elevated privileges can modify advanced power plan settings. The advanced plan settings are shown in Figure 11-29.




FIGURE 11-29 Advanced power settings

Using the Advanced Settings dialog box in Power Options, you can configure the following settings for both the On Battery and Plugged In states:

- **Require A Password On Wakeup** Specifies whether a password is required after the computer resumes from hibernation or sleep.
- **Turn Off Hard Disk** Specifies the period of inactivity after which to turn off the hard disk drive.
- **Desktop Background Settings** Specifies whether animated desktop backgrounds, such as the slideshow background, are available when the computer is on battery or plugged in.
- **Wireless Adapter Settings** Configure the power performance of the wireless adapter. Possible settings are Maximum Performance, Low Power Saving, Medium Power Saving, and Maximum Power Saving. Different settings influence the performance of the adapter, reducing the adapter's maximum range and speed.

Local Intranet sites dialog box, as shown in Figure 12-21. The default security level of this zone is Medium-Low. Protected Mode is not enabled by default for sites in this zone.




FIGURE 12-21 Local Intranet

- **Trusted Sites** Sites in the Trusted Sites zone often require elevated privileges. Only select Web sites should be added to this zone because Internet Explorer uses fewer security precautions with sites that are placed in this zone. The default security level for this zone is Medium. Protected Mode is not enabled by default for sites in this zone. Sites are added through the Trusted Sites dialog box. The default setting requires all sites in the Trusted Sites zone to be secured with a Secure Sockets Layer (SSL) certificate.
- **Restricted Sites** Sites in this zone are potentially malicious. You should only add sites to this zone if it is necessary to visit dangerous Web sites. The default security level for this zone is High. Internet Explorer Protected Mode is enabled by default for sites in this zone.
- **Internet** This zone hosts all Web sites that are not contained in the Local Intranet, Trusted Sites, or Restricted Sites zone. Sites in this zone are blocked from viewing private data from other Web sites. Sites in this zone are unable to make changes to Windows 7. The default security level of this zone is Medium-High. Protected Mode is enabled by default for sites in this zone.

You can use the slider, shown in Figure 12-22, to adjust the security level assigned to a zone. You can also configure whether a zone uses Protected Mode and Configure Custom Zone settings. Protected Mode is a technology that forces Internet Explorer to run as a low-integrity process. The security architecture of Windows 7 means that processes that are assigned lower integrities are unable to interact directly with objects that are assigned higher integrities. This means that any malware that might compromise the browser is blocked from causing damage to Windows 7 because it is unable to cause problems as a low-integrity process. The design of Windows 7 allows the processes that run in each tab to be separate from each other. This means that a tab that has a Web site in Protected Mode can run alongside a tab that has a site that is not running in Protected Mode. Sites that you do not trust, such as those on the Internet or within the Restricted Sites zone, are run in Protected Mode.



Quick Check

1. On which tab of the Performance Monitor Properties dialog box can you specify how often the graphs update?
2. Which rights does a user need to be able to monitor performance data remotely?

Quick Check Answers

1. On the General tab, in the Graph Elements group, you can adjust the Sample Every box to change how frequently the graph updates.
2. At a minimum, the user's account must be a member of the Performance Log Users group and the Event Log Readers group on the remote computer.

Data Collector Sets

Data collector sets (DCSs) gather system information, including configuration settings and performance data, and store it in a data file. You can use Performance Monitor to examine the data file and analyze detailed performance data, or you can generate a report that summarizes this information.

Windows 7 includes the following built-in DCSs:

- **System Performance** You can use this DCS when troubleshooting a slow computer or intermittent performance problems. It logs processor, disk, memory, and network performance (Internet Protocol versions 4 and 6) counters and kernel trace data.
- **System Diagnostics** You can use this DCS when troubleshooting reliability problems such as problematic hardware, driver failures, or STOP errors. It logs all the information included in the System Performance DCS, plus detailed system information. Figure 13-3 shows some of the counters included in the System Diagnostics data set.

To use a DCS, right-click it and then select Start. The System Performance DCS has a default overall duration of 1 minute. The System Diagnostics DCS collector set has a default overall duration of 10 minutes. To stop a DCS manually, right-click it and then click Stop.

After running a DCS, you can view a summary of the data that it has gathered in the *Performance Monitor\Reports* node. To view the most recent report for a DCS, right-click the DCS and then click *Latest Report*. You can then view the report by accessing it in the *Reports* node, as shown in Figure 13-4.

You can also add performance counter alerts to DCSs. This enables you to monitor and detect an alert, which you can then use to start a batch file, send you an e-mail, or call you on a pager. For example, if you configured an alert to trigger when free space on a logical volume falls below 30 percent, you could add this to a DCS and use it to trigger a batch file that archives the data on the volume.




FIGURE 13-13 The Performance tab in Task Manager

The next two graphs display how much RAM is being used, both at the moment and for the past few minutes. The percentage of memory being used is listed at the bottom of the Task Manager window. If memory use appears to be consistently high or slows your computer's performance noticeably, try reducing the number of programs that are open at one time (or encourage users you support to close any applications they are not currently using). If the problem persists, you might need to install more RAM or implement ReadyBoost.

Three tables below the graphs list various details about memory and resource usage. In the Physical Memory (MB) table, *Total* is the amount of RAM installed on your computer, *Cached* refers to the amount of physical memory used recently for system resources, and *Free* is the amount of memory that is currently unused and available.

In the Kernel Memory (MB) table, *Paged* refers to the amount of virtual memory the kernel is using; *Nonpaged* is the amount of RAM memory used by the kernel.

The System table has five fields: Handles, Threads, Processes, Up Time, and Page File. Handles are pointers that refer to system elements. They include (but are not limited to) files, registry keys, events, or directories. Lesson 2, "Configuring Performance Settings," discusses page file configuration.

If you need more information about how memory and CPU resources are being used, click Resource Monitor. This displays the Resource Monitor, which is discussed later in this lesson. You require elevated privileges to access Resource Monitor.

You can determine how much memory an individual process uses by selecting the Task Manager Processes tab. As shown in Figure 13-14, the Memory (Private Working Set) column is selected by default. A private working set indicates the amount of memory a process is using that other processes cannot share. This information can be useful in identifying

Lesson 2: Configuring Performance Settings

This lesson looks at configurations that can affect the performance of your computer and the tools that Windows 7 provides to display and reconfigure performance settings and resolve performance issues. If you do not like the tools provided, you can use Windows Management Instrumentation (WMI) scripts to write your own.

Many factors affect performance, such as the appearance of your screen or your browser window, the services and processes that are running on your computer, and the priorities and processor affinity that you assign to various processes. Performance is affected by your cache and page file settings, by the services and applications that start automatically or run even when not required, and by what processes are running and the amount of resources each consumes.

After this lesson you will be able to:

- Use a variety of Windows tools to inspect and configure settings that affect Windows 7 performance.
- Write WMI scripts that return system information and use the WMI tools.
- Troubleshoot performance issues.

Estimated lesson time: 45 minutes

Obtaining System Information Using WMI

WMI lets you access system management information and is designed to work across networks. It provides a consistent model of the managed environment and a WMI class for each manageable resource. A WMI class is a description of the properties of a managed resource and the actions that WMI can perform to manage that resource. A managed resource is any object (computer hardware, computer software, service, or user account) that can be managed by using WMI.

To use WMI, you write scripts that use the WMI scripting library. This library lets you work with WMI classes that correspond to managed resources. You can use this approach to manage resources such as disk drives, event logs, and installed software.

You can use Windows Script Host (WSH), Microsoft Visual Basic Scripting Edition (VBScript), Microsoft JScript, or scripting languages such as ActivePerl to write WMI scripts that automate the management of aspects of your network. Typically, Windows Management Instrumentation (WMI) files have `.vbs` extensions.

You can write scripts to manage event logs, file systems, printers, processes, registry settings, scheduled tasks, security, services, shared folders, and so on. You can create WMI-based scripts to manage network services, such as the Domain Name System (DNS), and to manage client-side network settings, such as whether a computer is configured with static

Internet Protocol version 4 (IPv4) address settings or whether it obtains these settings from a Dynamic Host Configuration Protocol (DHCP) server. WMI scripts can monitor and respond to entries in an event log, modifications to the file system or the registry, and other real-time operating system changes.

A WMI script works with WMI classes, which are representations of physical features or services on a computer. Each class can contain one or more objects or instances and the objects have attributes. You can display the value of each attribute or pass this on to another routine for analysis.

Typically, you type WMI scripts using a text editor such as Microsoft Notepad and save them as *.vbs* files in a directory (for example, *C:\WMI_Scripts*) that you have created for this purpose. Be wary of using word processing software such as Microsoft Office Word for this process. Word processing software often uses different styles of quotation marks for different fonts (to cite one example), and this can cause syntax errors. You can run WMI scripts from an elevated command prompt by using the *Cscript* utility, and you can create batch files that run scripts at scheduled intervals or when triggered by an event.

For example, the following WMI script accesses instances of the *Win32_Battery* class (there is only one) and prints out the value of the *EstimatedChargeRemaining* attribute. The code looks more complex than it actually is. You can substitute other WMI classes and find the values of their attributes by substituting the class and the attributes in this routine.

```
strComputer = "."
Set objSWbemServices = GetObject("winmgmts:\\." & strComputer)
Set colSWbemObjectSet = objSWbemServices.InstancesOf("Win32_Battery")
For Each objSWbemObject In colSWbemObjectSet
Wscript.Echo "Remaining Charge: " & objSWbemObject.EstimatedChargeRemaining & "
percent."
Next
```

Figure 13-29 shows the output from this script file, saved as *Battery.vbs* in the *C:\WMI_Scripts* folder. Note that if you run this script on a desktop computer, it should complete without error, but it does not give an output.




FIGURE 13-29 Estimated battery charge remaining read by a WMI script




FIGURE 13-40 The Boot tab of System Configuration

On the Startup tab, you can disable automatic startup for an application by clearing the check box beside the item. You can disable automatic startup for all items by clicking Disable All. This does not prevent the software from running—it merely stops it from starting automatically when the computer boots. The Services tab works in much the same way, in that you can disable or enable automatic startup of a single service or of all services. You can also determine what third-party services are running by selecting the Hide All Microsoft Services check box.

The Tools tab performs a very useful function. Not only are all the available tools listed, but you can enable any tool from this tab. This is often easier than trying to remember or deduce the tool's place in the Control Panel hierarchy, whether the tool is a Microsoft Management Console (MMC) snap-in, or what file you need to access from the command prompt to start the tool. The tab also lists the file and file path for the application that runs each tool.



EXAM TIP

You can use either Task Manager or the Services Console to start and stop services on a computer running Windows 7 without rebooting the computer.

Using the Services Console

The Services console, an MMC snap-in, lists the same services as does the Services tab of the System Configuration tool, but it provides more information about each service and more service management options. For example, the Services console tells you the service startup type (not just whether or not it is running) and the logon details.

You can access the Services console by entering **services.msc** in the Search box on the Start menu, in the Run box, or in a command-prompt window.

When you right-click a service in the Services console, you can start it, stop it, restart it, pause it, and resume it. You can access the Properties dialog box for the service and select the General, Log On, Recovery, and Dependencies tabs.

The General tab lets you specify the startup type. This can be Automatic, Automatic (Delayed Start) Manual, or Disabled. You should consider the following when specifying the startup type:

- If a service is configured as Automatic, it starts at boot time. Some services also automatically stop when no longer required. However, if you find that you do not need a service, configure its start type as Manual or Disabled.
- If a service is configured as Automatic (Delayed Start), it starts just after boot time. Configuring this setting can result in a faster boot, but if you need the service to be up and running when you boot, configure it as Automatic. If, on the other hand, you do not need a service, configure its start type as Manual or Disabled.
- Manual mode allows Windows 7 to start a service when needed. In practice, some services do not start up when required in Manual mode. If you find that you need a service, configure it as Automatic.
- If you configure a service as Disabled, it does not start even if needed. Unless you have a very good reason for disabling a service, configure its startup type as Manual instead.

The General tab, shown in Figure 13-41, also tells you whether a service is currently started, lets you start or stop it (as appropriate), and specifies the start parameters.




FIGURE 13-41 The General tab of the Service Properties dialog box

1. What WMI tool do you use to view Windows Management–generated events and event information, such as the event’s date and time, class, point of origin, and description?
 - A. WMI CIM Studio
 - B. WMI Object Browser
 - C. WMI Event Registration Tool
 - D. WMI Event Viewer
2. Which Windows Performance Analysis tool captures user and kernel traces and can merge them to form a combined trace?
 - A. Performance Analyzer
 - B. On/Off Transition Trace Capture
 - C. Trace Capture, Processing, and Command-Line Analysis
 - D. Visual Trace Analysis
3. Which tool provided by Windows 7 helps you determine which applications are responsible for activity on your hard disk, including which files and folders are being accessed?
 - A. Process Explorer
 - B. Resource Monitor
 - C. Task Manager
 - D. Windows Experience Index
4. A number of processor-intensive applications have been performing slowly on your computer. As a result, you add a second processor. This does not solve your problem, however, and you examine processor usage with Task Manager and Performance Monitor. You deduce that several key processes are using only the original processor. How do you ensure that these processes use whatever processor is available?
 - A. Configure Process Affinity on the Processes tab of Task Manager.
 - B. Configure Process Priority on the Processes tab of Task Manager.
 - C. Select Adjust For Best Performance Of Programs on the Advanced tab of the Performance Options tool.
 - D. Reconfigure Virtual Memory settings on the Advanced tab of the Performance Options tool.
5. Your computer is configured to dual-boot between Windows Vista Professional and Windows 7 Enterprise. Currently, it boots into Windows Vista by default. You want to specify Windows 7 as the startup default operating system and configure how Windows 7 reacts in the event of a system failure. You boot the computer into Windows 7. What tool do you use to accomplish your goal?
 - A. The Services console
 - B. Performance Options
 - C. Task Manager
 - D. System Configuration

✓ Quick Check

- All the client computers on your production network run Windows 7 Enterprise. They all have a single internal hard disk. You do not intend to provide an external hard disk for every client computer. You want to perform regular System Image backups. What type of backup destination would you use?

Quick Check Answer

- In this scenario, you would back up to a network share on either a storage network system or a file server.

Backup And Restore in Windows 7 supports backing up data files to CD-ROM, DVD-ROM, hard disk (including VHD files), or a network location. You can use Backup And Restore to write a System Image backup to an internal hard disk drive, an external hard disk drive (if formatted with the NTFS file system) and a network location. You cannot use Backup And Restore to write a System Image backup to a USB flash drive or a tape drive.

Bear in mind that you can save your backups on a network location only on computers running Windows 7 Professional, Windows 7 Ultimate, and Windows 7 Enterprise. Remember also that tape drives are not supported by the Backup And Restore utility.

NOTE BitLocker

You cannot select a backup destination that has BitLocker enabled.

File and Folder Backups

The Backup And Restore console graphical user interface (GUI) manually initiates backup and restore sessions and schedules automatic backups. You need to schedule client computers that store important data for automatic backup. After you first configure automatic file backup using the Backup And Restore console, Windows 7 regularly backs up your files. The first time a backup is performed, a full backup is done, including all important user documents. Subsequent backups are incremental, backing up only changed files. Older backups are discarded if the disk begins to run out of space.

For example, if you configure a nightly scheduled backup and change a file every day, a copy of that file is stored in each day's Backup Files folder. By storing multiple versions of a single file, Windows 7 gives users the opportunity to choose from several older copies of a file when using the Previous Versions tool (described in Lesson 3, "Recovering Files and Folders"). When you restore files, you need only restore from a single backup because Windows 7 automatically locates the most recent version of each file. Windows 7 uses shadow copy (described in Lesson 3) to back up the last saved version of a file. Therefore, if a file is open during the backup, that file will be backed up. However, any changes the user made since last saving the file are not backed up.

You need administrator credentials to configure scheduled backups or to manually initiate a backup. However, restoring files does not require administrator privileges unless a user attempts to restore another user's file.

If you perform a file backup to a shared network folder, the credentials used to run the backup must have Full Control share and NTFS permissions for the destination folder (known as Co-owner permissions in the Windows 7 Setup Wizard). To reduce security risks, you should set up a user account that is used only by the backup application and configure share and NTFS permissions to grant access only to the backup user. The backup account requires administrative privileges to the computer being backed up, but it needs permissions only to the share and folder on the target computer.

File and Folder Backup Structure

When a user chooses to perform a backup to an external hard disk, Windows 7 automatically creates a folder in the root of the hard disk using the computer name. Within that folder, backups are saved in the format: *Backup Set <year-month-day> <time>*. For example, if your computer name is Canberra, your backup location is H:, and you backed up on July 4, 2009, at 16:36:39, your backup is located in *H:\Canberra\Backup Set 2009-07-04 163639*, as shown in Figure 14-4.




FIGURE 14-4 Backup set

Lesson 2: System Recovery

This lesson discusses how to restore corrupt or misconfigured system settings to a previous working system configuration. It looks at system protection, restore points, system recovery, and restoring a System Image.

The chapter considers boot options, advanced recovery methods, and how you recover from the situation when a system change prevents your computer from booting normally.

After this lesson you will be able to:

- Create a restore point manually.
- Perform a system restore to a selected restore point.
- Restore from a System Image backup.
- Boot from the Windows 7 installation DVD-ROM and run a system repair.
- Use the Advanced Boot Options.
- Configure System Protection.

Estimated lesson time: 30 minutes

Performing a System Restore

Windows 7 creates system restore points on a regular schedule and prior to events such as the installation of applications and drivers. A *restore point* contains information about registry settings and other system information. Windows 7 generates restore points automatically before implementing significant system changes. You can manually create restore points and restore a computer system to a selected restore point.

If you install an application or driver that causes your computer to become unstable, you should first attempt to uninstall the application or roll back the driver. If this does not solve the problem, you can restore system files and settings by performing a system restore to restore the computer to its last system restore point. A *system restore* returns a computer system to a selected restore point. System restores do not alter user files. Note that a system restore is not the same as a System Image restore, which is discussed later in this lesson.

The most straightforward way of starting a system restore is to enter **system restore** in the Start menu Search box and click the appropriate link. In Control Panel, you can also access All Control Panel Items, click Recovery, and click Open System Restore. In addition, you can start a system restore from the System Protection tab of the System Properties dialog box. You do this in the Practice Session later in this lesson. If you boot from the installation media, you can select Repair Your Computer and then select Use The System Restore option. You need elevated privileges to run system restore. The System Restore Wizard is shown in Figure 14-10.




FIGURE 14-11 The Advanced Recovery Methods dialog box

The first method uses a System Image backup, which was described in Lesson 1. The second method reinstalls Windows, either from a recovery image provided by your computer manufacturer or from the original Windows installation files. Both methods can result in loss of data. Before beginning either method, you are prompted to back up your personal files to an external location, such as a USB hard disk. The procedure for restoring from a System Image backup is described in the next section. Note that restoring from a System Image backup is not the same as performing a system restore.

Restoring from a System Image Backup

A System Image restore rewrites the entire contents of a system volume. Therefore, you restore from a System Image backup by booting from the Windows 7 Installation DVD-ROM and loading System Recovery tools or by pressing F8 during the boot process. Restoring from a System Image backup enables you to quickly get a computer running after you replace a failed hard disk, or if the operating system installation has been corrupted (for example, by malware that cannot be removed except by wiping the disk). It is sometimes known as *complete recovery* or *complete PC Restore*.

This procedure assumes that the System Recovery Options (otherwise known as the Windows Recovery Environment, or Windows RE) files are present on the DVD-ROM. If not, you can boot from the installation DVD-ROM and press F8 during the boot to access the Advanced Boot Options, as described in the next section of this lesson.

3. Correct Answers: B and D

- A. Incorrect:** An x86 version of Windows 7 can utilize a maximum of 4 GB of RAM.
- B. Correct:** An x64 versions of Windows 7 can utilize more than 4 GB of RAM.
- C. Incorrect:** An x86 version of Windows 7 can utilize a maximum of 4 GB of RAM.
- D. Correct:** An x64 versions of Windows 7 can utilize more than 4 GB of RAM.

4. Correct Answer: B

- A. Incorrect:** A computer does not require a DVD-ROM drive for a network-based installation from WDS.
- B. Correct:** To deploy over the network using WDS requires the computer either a PXE-compliant network adapter or be booted from a WDS discover image.
- C. Incorrect:** A computer does not require a USB 2.0 slot for a network-based installation from WDS.
- D. Incorrect:** A computer does not require a HDMI port for a network-based installation from WDS.

5. Correct Answer: B

- A. Incorrect:** Two volumes are the minimum number required to support dual-booting between Windows XP and Windows 7.
- B. Correct:** Two volumes are the minimum number required to support dual-booting between Windows XP and Windows 7. The Windows 7 installation routine creates an extra 200-MB system volume on one of these volumes when you install Windows 7, but two volumes is the minimum number required on a Windows XP computer prior to attempting to install Windows 7 for dual-boot configuration.
- C. Incorrect:** Two volumes are the minimum number required to support dual-booting between Windows XP and Windows 7.
- D. Incorrect:** Two volumes are the minimum number required to support dual-booting between Windows XP and Windows 7.

Lesson 2

1. Correct Answer: B

- A. Incorrect:** You cannot upgrade Windows Vista Enterprise to Windows 7 Home Professional. Windows Vista Enterprise (x86) can only be upgraded to Windows 7 Enterprise (x86).
- B. Correct:** Windows Vista Enterprise (x86) can only be upgraded to Windows 7 Enterprise (x86).
- C. Incorrect:** You cannot upgrade an x86 version of Windows Vista to an x64 version of Windows 7.
- D. Incorrect:** You cannot upgrade an x86 version of Windows Vista to an x64 version of Windows 7.

Case Scenario 2: Working with VHDs

1. You can boot only the computer running Windows 7 from the operating system captured on its VHD. Only Windows 7 Enterprise and Windows 7 Ultimate can be booted when installed on a VHD.
2. You need to run *sysprep /generalize* to strip hardware-specific information such as the SID from the image generated from the reference computer.

Chapter 3: Lesson Review Answers

Lesson 1

1. **Correct Answers: B and D**
 - A. Incorrect:** The *DISM /image* option specifies the folder that contains the mounted image, not the source image.
 - B. Correct:** The */recurse* parameter causes all the drivers in the folder *C:\drivers* and its subfolders to be added to the mounted image.
 - C. Incorrect:** The *DISM /image* option specifies the folder that contains the mounted image, not the source image.
 - D. Correct:** This command would work, although Answer B would be the preferred solution. Typically, you would use the */driver* parameter multiple times if the drivers were not in the same file structure (for example, *C:\Printerdriver* and *C:\Scannerdriver*).
2. **Correct Answer: A**
 - A. Correct:** The amount of writeable space available on a Windows PE system volume when booted in RAMdisk mode is known as the Windows PE system volume scratch space. You can use the */get-scratchspace* DISM option to obtain a value for this.
 - B. Incorrect:** This command returns the path to the root of the Windows PE image at boot time, known as the target path.
 - C. Incorrect:** This command determines whether the Windows PE profiling tool is enabled or disabled.
 - D. Incorrect:** Profiling (file logging) is disabled by default. This command enables it. However, the command does not return the amount of writeable space available on a Windows PE system volume when booted in RAMdisk mode.
3. **Correct Answer: D**
 - A. Incorrect:** The */set-syslocale* option sets the language for non-Unicode programs (also called the system locale) and the font settings. You can do this only in an offline-mounted image.
 - B. Incorrect:** The */set-userlocale* option configures a per-user setting that determines the default sort order and the default settings for formatting dates, times, currency, and numbers. You can do this only in an offline-mounted image.

Chapter 5: Case Scenario Answers

Case Scenario 1: Configuring Application Compatibility at Fabrikam

1. Edit the properties of application Alpha. Configure the application to run using the Windows XP Service Pack 3 compatibility mode.
2. Edit the properties of the Beta application. On the compatibility tab, enable the Run This Program As An Administrator option. This enables the Run As Administrator option without having to right-click the application each time to enable this functionality.
3. You can use the ACT to configure compatibility options for Application Gamma.

Case Scenario 2: Restricting Applications at Contoso

1. Configure an AppLocker executable rule that uses a file hash of the data collection application. You cannot use a publisher rule because the application is not digitally signed.
2. Configure this rule to apply to the Scientists group. Allow the execution of the application. Ensure that the data collection application is installed outside the scope of the folders covered by the default allow rule so that it will be blocked unless specifically allowed.
3. The in-house developers would need to sign the application digitally before you can create a publisher rule for it.

Chapter 6: Lesson Review Answers

Lesson 1

1. **Correct Answer: C**
 - A. Incorrect:** The *Ping* command is used to test connectivity. It does not display the IP configuration of a computer's interfaces.
 - B. Incorrect:** The *Tracert* command is used to test connectivity to a device on a remote network and return information about the intermediate hops. It does not display the IP configuration of a computer's interfaces.
 - C. Correct:** The *Ipconfig* command displays the IP configuration of a computer's interfaces.
 - D. Incorrect:** The *Netstat* tool displays protocol statistics. It does not display the IP configuration of a computer's interfaces.
2. **Correct Answers: B, C, and D**
 - A. Incorrect:** This accesses the Local Area Network (LAN Settings) dialog box. You can select automatic configuration, specify an automatic configuration script, or specify a proxy server. The dialog box does not display connection properties.

4. Correct Answer: B

- A. Incorrect:** Although you can create rules based on applications using Windows Firewall, you cannot use this tool to create rules that require that incoming connections be authenticated.
- B. Correct:** WFAS allows you to create detailed rules that include the ability to allow incoming traffic only if it is authenticated.
- C. Incorrect:** Credential Manager stores authentication credentials. It cannot be used to create firewall rules that require authentication.
- D. Incorrect:** Authorization Manager allows you to configure roles for the delegation of administrative privileges. You cannot use Authorization Manager to create firewall rules that require authentication.

5. Correct Answers: A and D

- A. Correct:** You should configure Windows Firewall to notify you when it blocks a program in the Home Or Work (Private) Network Location Settings area. This ensures that you receive a message when a new program is blocked when connected to this network profile.
- B. Incorrect:** You should not disable the setting related to receiving a message when a new program is blocked in the Home Or Work (Private) Network Location Settings area because this means that you do not receive a message when a program is blocked.
- C. Incorrect:** You should not enable the setting related to receiving a message when a new program is blocked in the Public Network Location Settings area because this notifies you when a new program is blocked. The question text states that you should not be notified when this occurs.
- D. Correct:** You should disable the setting related to receiving a message when a new program is blocked in the Public Network Location Settings area because this ensures that you are not notified when a program is blocked.

Lesson 2

1. Correct Answer: C

- A. Incorrect:** You should not enable Remote Assistance. Remote Assistance requires that someone is logged on to the computer that you wish to manage remotely.
- B. Incorrect:** You should not enable the Remote Desktop: Don't Allow Connections To This Computer option because that blocks the ability to make Remote Desktop connections.
- C. Correct:** You should enable the Remote Desktop: Allow Connections From Computer Running Any Version Of Remote Desktop setting because this allows you to connect to a computer running Windows 7 from a computer running Windows XP with SP2.
- D. Incorrect:** You should not enable the Remote Desktop: Allow Connections Only From Computers With Network Level Authentication as clients running Windows XP with SP2 are unable to connect to clients running Windows 7 when this option is enabled. Windows XP requires SP3 and special configuration to use Network-Level Authentication.

Chapter 13: Lesson Review Answers

Lesson 1

1. **Correct Answer: A**
 - A. Correct:** The lowest subscore determines the base score, even though this computer is not primarily used for three-dimensional graphics and gaming.
 - B. Incorrect:** 3.9 is the average of the subscores. However, the lowest subscore determines the base score, not the average.
 - C. Incorrect:** 5.1 is a significant score because the computer is used for processor-intensive operations. However, the question asks for the base score, and the lowest subscore determines the base score.
 - D. Incorrect:** 5.3 is the highest subscore. The lowest subscore determines the base score, not the highest.
2. **Correct Answer: A**
 - A. Correct:** Reliability Monitor tracks application installations. It enables you to determine whether what applications have been installed and exactly when the installations occurred.
 - B. Incorrect:** Action Center can tell you if an application or device driver is not working properly. It cannot tell you when an application was installed.
 - C. Incorrect:** DCSs capture current performance and configuration data. They cannot tell you when in the past an application was installed.
 - D. Incorrect:** You can use Performance Monitor to view performance counters in real time or analyze performance data in a DCS. However, Performance Monitor does not record when an application was installed.
3. **Correct Answers: A, B, C, and E**
 - A. Correct:** Application failures are recorded in Reliability Monitor.
 - B. Correct:** Windows errors are recorded in Reliability Monitor.
 - C. Correct:** Application installs and uninstalls are recorded in Reliability Monitor.
 - D. Incorrect:** A service starting or stopping is typically recorded in the event log, not Reliability Monitor.
 - E. Correct:** Device driver failures are recorded in Reliability Monitor.
4. **Correct Answer: A**
 - A. Correct:** You can use the Wecutil utility to configure the Event Collector service.
 - B. Incorrect:** The *Winrm* command configures WinRM. Typically, you run it on a source computer. You can run it on the collector computer if you are configuring a source-initiated conscription, but this is not relevant to this scenario because Canberra is retrieving events from Aberdeen. In any case, this command does not configure the Event Collector service.
 - C. Incorrect:** The *Winrm* command configures WinRM. Here, you are running it in quiet mode. Whether you use the *-q* switch or not, the command does not configure the Event Collector service.

3. Correct Answers: A, B, C, and D

- A. Correct:** Don's Word and Excel files in his Documents library are backed up by default. Don does not need to be logged into the computer.
- B. Correct:** Kim's Word and Excel files in her Documents library are backed up by default.
- C. Correct:** Don's Presentation folder is in his Documents library and is backed up by default.
- D. Correct:** Don's Picture library is backed up by default.
- E. Incorrect:** Encrypted files and folders are not backed up (even if they are in a Documents library).
- F. Incorrect:** Files in the Recycle Bin are not backed up.
- G. Incorrect:** By default, a USB flash memory device is formatted with the FAT filing system. Files on such a device are not backed up.

4. Correct Answers: A, B, D, and E.

- A. Correct:** You can use Backup And Restore to write a System Image backup to an internal hard disk.
- B. Correct:** You can use Backup And Restore to write a System Image backup to an external hard disk (if formatted with the NTFS file system).
- C. Incorrect:** You cannot use Backup And Restore to write a System Image backup to an 8-GB USB flash drive, although you can use this media for file and folder backups.
- D. Correct:** You can use Backup And Restore to write a System Image backup to an optical drive.
- E. Correct:** You can use Backup And Restore to write a System Image backup to a network share provided the computer is running **Windows 7 Professional**, **Windows 7 Ultimate**, or **Windows 7 Enterprise**.

5. Correct Answer: D

- A. Incorrect:** Computers running Windows 7 Ultimate offer all Windows 7 features. However, they are not the only computers running Windows 7 that can be backed up to a network share.
- B. Incorrect:** Computers running Windows 7 Enterprise are installed with a bulk license and are intended for use in large numbers in the enterprise environment. However, they are not the only computers running Windows 7 that can be backed up to a network share.
- C. Incorrect:** Some features, such as booting from a VHD, require either Windows 7 Ultimate or Windows 7 Enterprise. However, the ability to back up to a network share is not one of these features.
- D. Correct:** You can save backups to a network share on computers running **Windows 7 Professional**, **Windows 7 Ultimate**, and **Windows 7 Enterprise**.