Part I
Learn at Your Own Pace

1 Introducing Microsoft Windows Server 2003

This chapter does not cover specific exam objectives. After introducing the Microsoft Windows Server 2003 family of products, this chapter covers some installation and configuration considerations with a focus on what you need to know for the 70-290 certification exam.

Why This Chapter Matters

The purpose of this book is to empower you to manage and maintain a Microsoft Windows Server 2003 environment, and to prepare you effectively for the 70-290 certification examination. Although it is assumed that you have experience with Microsoft Windows technologies, the Windows Server 2003 family and Microsoft Active Directory directory service itself may be new to you. The goal of this chapter, therefore, is to introduce you to the multiple versions and editions of Windows Server 2003, so that you can identify the key distinctions among them and determine the mix of versions that will most effectively meet the needs of your organization. You will then be guided through the process of installing and configuring a Windows Server 2003 computer that functions as a domain controller in an Active Directory domain.

Lessons in this Chapter:

Before You Begin

This chapter will guide you through the steps required to configure a computer running Windows Server 2003. You will be able to use that computer for the hands-on exercises throughout this training kit. The computer should have at least one disk drive that can be erased and used to install Windows Server 2003.

Lesson 1: The Windows Server 2003 Family

Windows Server 2003 is, of course, more secure, more reliable, more available, and easier to administer than any previous version of Windows. Let's take a close look at the platform and how it compares to Microsoft Windows 2000. This lesson provides a brief overview of the Windows Server 2003 family, focusing on the differences among the product editions: Web Edition, Standard Edition, Enterprise Edition, and Datacenter Edition.

After this lesson, you will be able to

- Identify the key differences among the Windows Server 2003 versions

Estimated lesson time: 5 minutes

Windows Server 2003 Editions

Windows Server 2003 is an incremental update to the platform and technologies introduced in Windows 2000. If you are coming to Windows Server 2003 with experience from Windows 2000 servers, you will find the transition a relatively easy one. If your experience is with Windows NT 4, welcome to the new world!

But don't let the incremental nature of the updates mislead you; behind the upgrades are significant and long-awaited improvements to the security and reliability of the operating system and to the administrative toolset. In many books, this would be the place where you would get a laundry list of new features. Actually, the Windows Server 2003 list is extensive and there are features that make upgrading to Windows Server 2003 an obvious choice for almost any administrator. However, the particular features that appeal to you may be different from those that appeal to another IT professional.

You may be drawn to the significant features and improvements added to Active Directory, the new tools to support popular but complex GPOs, the enhancements to enterprise security, the improvements to Terminal Services, or a number of other enhanced capabilities of the new operating system. If you are considering a move to Windows Server 2003, take a good look through the Microsoft Web site for the platform, at *http://www.microsoft.com/windowsserver2003* and judge for yourself which improvements are, in your environment, truly significant.

Although the list of new features is extensive, the evaluation of the operating system becomes more interesting because Windows Server 2003 is available in multiple flavors including the 32-bit, 64-bit, and embedded versions. But the most important distinc-

tions are those among the four product editions, listed here in order of available features and functionality, as well as by price:

- Windows Server 2003, Web Edition
- Windows Server 2003, Standard Edition
- Windows Server 2003, Enterprise Edition
- Windows Server 2003, Datacenter Edition

Web Edition

To position Windows Server 2003 more competitively against other Web servers, Microsoft has released a stripped-down-yet-impressive edition of Windows Server 2003 designed specifically for Web services. The feature set and licensing allows customers easy deployment of Web pages, Web sites, Web applications, and Web services.

Web Edition supports 2 gigabytes (GB) of RAM and a two-way symmetric multiprocessor (SMP). It provides unlimited anonymous Web connections but only 10 inbound server message block (SMB) connections, which should be more than enough for content publishing. The server cannot be an Internet gateway, DHCP or fax server. Although you can remotely administer the server with Remote Desktop, the server cannot be a terminal server in the traditional sense. The server can belong to a domain, but cannot be a domain controller. The included version of the Microsoft SQL Server Database Engine can support as many as 25 concurrent connections.

Standard Edition

Windows Server 2003, Standard Edition, is a robust, multipurpose server capable of providing directory, file, print, application, multimedia, and Web services for small to medium-sized businesses. Its comprehensive feature set is expanded, compared to Windows 2000, with Microsoft SQL Server Database Engine (MSDE), a version of SQL Server that supports five concurrent connections to databases up to 2 GB in size; a free, out-of-the-box Post Office Protocol version 3 (POP3) service which, combined with the included Simple Mail Transfer Protocol (SMTP) service, allows a server to function as a small, stand-alone mail server; and Network Load Balancing (NLB), a useful tool that was only included with the Advanced Server edition of Windows 2000.

The Standard Edition of Windows Server 2003 supports up to 4 GB of RAM and four-way SMP.

Note Through Release Candidate (RC) 1 of Windows Server 2003, the beta and prerelease code supported only two processors. That limitation was removed in RC2, and the Standard Edition supports four processors. Documentation and resources that were created prior to release may contain misleading information regarding SMP support.

Enterprise Edition

The Enterprise Edition of Windows Server 2003 is designed to be a powerful server platform for medium- to large-sized businesses. Its enterprise-class features include support for eight processors, 32 GB of RAM, eight-node clustering (including clustering based on a Storage Area Network (SAN) and geographically dispersed clustering) and availability for 64-bit Intel Itanium-based computers, on which scalability increases to 64 GB of RAM and 8-way SMP.

Other features that distinguish the Enterprise Edition from the Standard Edition include:

- Support for Microsoft Metadirectory Services (MMS), which enables the integration of multiple directories, databases, and files with Active Directory.
- Hot Add Memory, so that you can add memory to supported hardware systems without downtime or reboot.
- Windows System Resource Manager (WSRM), which supports the allocation of CPU and memory resources on a per-application basis.

Datacenter Edition

The Datacenter Edition, which is available only as an OEM version as part of a high-end server hardware package, provides almost unfathomable scalability, with support on 32-bit platforms for 32-way SMP with 64 GB of RAM and on 64-bit platforms for 64-way SMP with 512 GB of RAM. There is also a 128-way SMP version that supports two 64-way SMP partitions.

64-Bit Editions

The 64-bit editions of Windows Server 2003, which run on Intel Itanium-based computers, provide for higher CPU clock speeds and faster floating-point processor operations than the 32-bit editions of Windows. CPU coding improvements and processing enhancements yield significantly faster computational operations. Increased access speed to an enormous memory address space allows for smooth operation of complex, resource-intensive applications, such as massive database applications, scientific analysis applications, and heavily accessed Web servers.

Some features of the 32-bit editions are not available in the 64-bit editions. Most notably, the 64-bit editions do not support 16-bit Windows application, real-mode applications, POSIX applications, or print services for Apple Macintosh clients.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the

question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You are planning the deployment of Windows Server 2003 computers for a department of 250 employees. The server will host the home directories and shared folders for the department, and it will serve several printers to which departmental documents are sent. Which edition of Windows Server 2003 will provide the most cost-effective solution for the department?

2. You are planning the deployment of Windows Server 2003 computers for a new Active Directory domain in a large corporation that includes multiple separate Active Directories maintained by each of the corporation's subsidiaries. The company has decided to roll out Exchange Server 2003 as a unified messaging platform for all the subsidiaries, and plans to use Microsoft Metadirectory Services (MMS) to synchronize appropriate properties of objects throughout the organization. Which edition of Windows Server 2003 will provide the most cost-effective solution for this deployment?

3. You are rolling out servers to provide Internet access to your company's e-commerce application. You anticipate four servers dedicated to the front-end Web application and one server for a robust, active SQL database. Which editions will provide the most cost-effective solution?

Lesson Summary

- Windows Server 2003 is available in 64-bit as well as 32-bit versions.
- The primary distinctions among versions of Windows Server 2003 are the product editions: Web Edition, Standard Edition, Enterprise Edition, and Datacenter Edition, each of which supports a subset of features honed to a specific purpose.
- Taken as a whole, Windows Server 2003 is an upgrade to Windows 2000. However, the feature and security improvements are significant, and you are likely to find that particular upgrades provide critical enhancements for your particular environment.

Lesson 2: Installation and Configuration of Windows Server 2003 and Active Directory

The 70-290 examination focuses on the management and maintenance of a Windows Server 2003 environment. The objectives of the exam focus very little attention on Active Directory itself; some of the objectives, however, relate to the administration of Active Directory objects: users, groups, computers, printers, and shared folders in particular. The chapters that follow will explain the examination objectives in detail, and hands-on exercises will be an important component of your learning experience. Those exercises require you to have configured a domain controller running Windows Server 2003. If you are comfortable configuring a domain controller and creating basic user, group, and computer accounts, you can skip this lesson. If you are less familiar with Active Directory, this lesson will provide sufficient foundation for you to embark on a full exploration of Windows Server 2003.

After this lesson, you will be able to

- Install Windows Server 2003
- Identify the key structures and concepts of Active Directory
- Create a domain controller
- Create Active Directory objects including users, groups, and organizational units (OUs)

Estimated lesson time: 60 minutes

Installing and Configuring Windows Server 2003

As an experienced IT professional, you have no doubt spent considerable time installing Windows platforms. Some of the important and enhanced considerations when installing Windows Server 2003 are

- **Bootable CD-ROM installation** Most administrators first became accustomed to installing an operating system by booting from the CD-ROM in the late 1990s. Windows Server 2003 continues the trend, and can be installed directly from the CD-ROM. But Windows Server 2003 adds a twist: there is *no* support for starting installation from floppy disks.
- **Improved graphical user interface (GUI) during setup** Windows Server 2003 uses a GUI during setup that resembles that of Windows XP. It communicates more clearly the current state of the installation and the amount of time required to complete installation.
- **Product activation** Retail and evaluation versions of Windows Server 2003 require that you activate the product. Volume licensing programs, such as Open License, Select License, or Enterprise Agreement, do not require activation.

The specific steps required to install Windows Server 2003 are outlined in Exercise 1.

After installing and activating Windows, you can configure the server using a well-thought-out Manage Your Server page, as shown in Figure 1-1, that launches automatically at logon. The page facilitates the installation of specific services, tools, and configurations based on server roles. Click Add Or Remove A Role and the Configure Your Server Wizard appears.

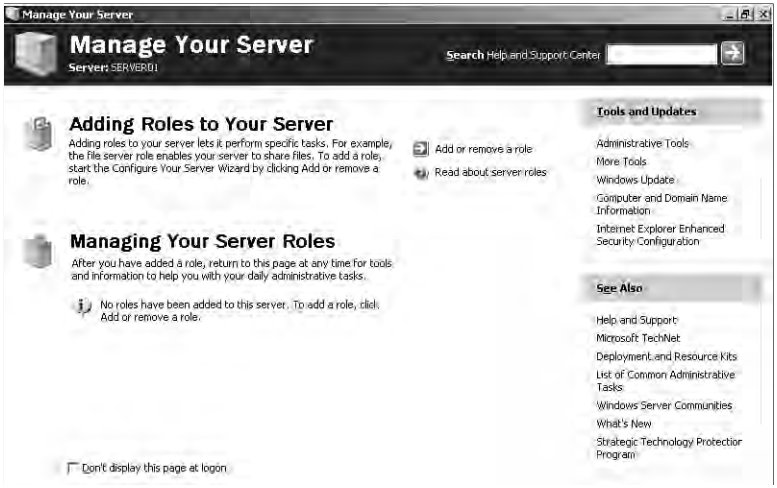

Figure 1-1 The Manage Your Server page

If you select Typical Configuration For A First Server, the Configure Your Server Wizard promotes the server to a domain controller in a new domain, installs Active Directory services, and, if needed, Domain Name Service (DNS), Dynamic Host Configuration Protocol (DHCP), and Routing And Remote Access (RRAS) service.

If you select Custom Configuration, the Configure Your Server Wizard can configure the following roles:

- **File Server** Provides convenient, centralized access to files and directories for individual users, departments, and entire organizations. Choosing this option allows you to manage user disk space by enabling and configuring disk quota management and to provide improved file system search performance by enabling the Indexing service.
- **Print Server** Provides centralized and managed access to printing devices by serving shared printers and printer drivers to client computers. Choosing this option starts the Add Printer Wizard to install printers and their associated Windows printer drivers. It also installs Internet Information Services (IIS 6.0) and configures Internet Printing Protocol (IPP) and installs the Web-based printer administration tools.

- **Application Server (IIS, ASP.NET)** Provides infrastructure components required to support the hosting of Web applications. This role installs and configures IIS 6.0 as well as ASP.NET and COM+.
- **Mail Server (POP3, SMTP)** Installs POP3 and SMTP so that the server can act as an e-mail server for POP3 clients.
- **Terminal Server** Provides applications and server resources, such as printers and storage, to multiple users as if those applications and resources were installed on their own computers. Users connect with the Terminal Services or Remote Desktop clients. Unlike Windows 2000, Windows Server 2003 provides Remote Desktop for Administration automatically. Terminal Server roles are required only when hosting applications for users on a terminal server.
- **Remote Access/VPN Server** Provides multiple-protocol routing and remote access services for dial-in, local area networks (LANs) and wide area networks (WANs). Virtual private network (VPN) connections allow remote sites and users to connect securely to the network using standard Internet connections.
- **Domain Controller (Active Directory)** Provides directory services to clients in the network. This option configures a domain controller for a new or existing domain and installs DNS. Choosing this option runs the Active Directory Installation Wizard.
- **DNS Server** Provides host name resolution by translating host names to IP addresses (forward lookups) and IP addresses to host names (reverse lookups). Choosing this option installs the DNS service, and then starts the Configure A DNS Server Wizard.
- **DHCP Server** Provides automatic IP addressing services to clients configured to use dynamic IP addressing. Choosing this option installs DHCP services and then starts the New Scope Wizard to define one or more IP address scopes in the network.
- **Streaming Media Server** Provides Windows Media Services (WMS). WMS enables the server to stream multimedia content over an intranet or the Internet. Content can be stored and delivered on demand or delivered in real time. Choosing this option installs WMS.
- **WINS Server** Provides computer name resolution by translating NetBIOS names to IP addresses. It is not necessary to install Windows Internet Name Service (WINS) unless you are supporting legacy operating systems, such as Windows 95 or Windows NT. Operating systems such as Windows 2000 and Windows XP do not require WINS, although legacy applications on those platforms may very well require NetBIOS name resolution. Choosing this option installs WINS.

To complete the hands-on exercises in this book, you will configure a computer as Server01, acting as a domain controller in the domain *contoso.com*. The steps for configuring the server as a domain controller using the Configure Your Server Wizard are listed in Exercise 2 at the end of this lesson.

Active Directory

Many books have been devoted to the planning, implementation, and support of Active Directory. If you're experienced with Active Directory, you will recognize that the following discussion has been simplified solely because it would take many books to discuss all the detail. The goal of this section is to distill that information to what you should know to approach the 70-290 exam.

Networks, Directory Services, and Domain Controllers

Networks were created on the day when the first user decided he or she didn't want to walk down the hall to get something from another user. In the end, networks are all about providing resources remotely. Those resources are often files, folders, and printers. Over time those resources have come to include many things, most significantly, e-mail, databases, and applications. There has to be some mechanism to keep track of these resources, providing, at a minimum, a directory of users and groups so that the resources can be secured against undesired access.

Microsoft Windows networks support two directory service models: the workgroup and the domain. The domain model is by far the more common in organizations implementing Windows Server 2003. The domain model is characterized by a single directory of enterprise resources—Active Directory—that is trusted by all secure systems that belong to the domain. Those systems can therefore use the security principals (user, group, and computer accounts) in the directory to secure their resources. Active Directory thus acts as an identity store, providing a single trusted list of Who's Who in the domain.

Active Directory itself is more than just a database, though. It is a collection of supporting files including transaction logs and the system volume, or Sysvol, that contains logon scripts and group policy information. It is the services that support and use the database, including Lightweight Directory Access Protocol (LDAP), Kerberos security protocol, replication processes, and the File Replication Service (FRS). The database and its services are installed on one or more domain controllers. A domain controller is a server that has been promoted by running the Active Directory Installation Wizard by running DCPROMO from the command line or, as you will do in Exercise 2, by running the Configure Your Server Wizard. Once a server has become a domain controller, it hosts a copy, or replica, of Active Directory and changes to the database on any domain controller are replicated to all domain controllers within the domain.

Domains, Trees and Forests

Active Directory cannot exist without at least one domain, and vice versa. A domain is the core administrative unit of the Windows Server 2003 directory service. However, an enterprise may have more than one domain in its Active Directory. Multiple domain models create logical structures called *trees* when they share contiguous DNS names. For example *contoso.com*, *us.contoso.com*, and *europe.contoso.com* share contiguous DNS namespace, and would therefore be referred to as a tree.

If domains in an Active Directory do not share a common root domain, they create multiple trees. That leads you to the largest structure in an Active Directory: the *forest*. An Active Directory forest includes all domains within that Active Directory. A forest may contain multiple domains in multiple trees, or just one domain. When more than one domain exists, a component of Active Directory called the Global Catalog becomes important because it provides information about objects that are located in other domains in the forest.

Objects and Organizational Units (OUs)

Enterprise resources are represented in Active Directory as objects, or records in the database. Each object has numerous attributes, or properties, that define it. For example, a user object includes the user name and password; a group object includes the group name and a list of its members.

To create an object in Active Directory, open the Active Directory Users And Computers console from the Administrative Tools program group. Expand the domain to reveal its containers and OUs. Right-click a container or OU and select New *object_type*.

Active Directory is capable of hosting millions of objects, including users, groups, computers, printers, shared folders, sites, site links, Group Policy Objects (GPOs), and even DNS zones and host records. You can imagine that without some kind of structure, accessing and administering the directory would be a nightmare.

Structure is the function of a specific object type called an organizational unit, or OU. OUs are containers within a domain that allow you to group objects that share common administration or configuration. But they do more than just organize Active Directory objects. They provide important administrative capabilities, as they provide a point at which administrative functions can be delegated and to which group policies can be linked.

Delegation

Administrative delegation relates to the simple idea that you might want a front-line administrator to be able to change the password for a certain subset of users. Each

object in Active Directory (in this case, the user objects) includes an access control list (ACL) that defines permissions for that object, just as files on a disk volume have ACLs that define access for those files. So, for example, a user object's ACL will define what groups are allowed to reset its password. It would get complicated to assign the front-line administrator permissions to change each individual user's password, so instead you can put all of those users in a single OU and assign that administrator the reset password permission on the OU. That permission will be inherited by all user objects in the OU, thereby allowing that administrator to modify permissions for all users.

Resetting user passwords is just one example of administrative delegation. There are thousands of combinations of permissions that could be assigned to groups administering and supporting Active Directory. OUs allow an enterprise to create an active representation of its administrative model, and to specify who can do what to objects in the domain.

Group Policy

OUs are also used to collect objects—computers and users—that are configured similarly. Just about any configuration you can make to a system can be managed centrally through a feature of Active Directory called Group Policy. Group Policy allows you to specify security settings, deploy software, and configure operating system and application behavior without ever touching a machine. You simply implement your configuration within a GPO.

GPOs are collections of hundreds of possible configuration settings, from user logon rights and privileges to the software that is allowed to be run on a system. A GPO is linked to a container within Active Directory—typically to an OU, but can also be domains, or even sites—and all the users and computers beneath that container are affected by the settings contained in the GPO.

You will likely see Group Policy referred to on the 70-290 exam. The important things to remember about Group Policy are that it is a tool that can centrally implement configuration; that some settings apply to computers only and some settings apply to users only; and that the only computers or users that will be affected by a policy are those that are beneath the OU to which the policy is linked.

Learning More

As suggested earlier in this section, Active Directory is a large and complex topic that deserves significant examination if you are going to implement Windows Server 2003 as a domain controller. The following Microsoft Press titles are recommended reading:

- *Active Directory for Microsoft Windows Server 2003 Technical Reference*
- *MCSE Self-Paced Training Kit (Exam 70-294): Planning, Implementing, and Maintaining a Microsoft Windows Server 2003 Active Directory Infrastructure*

Practice: Installing Windows Server 2003

In this practice, you will configure a computer to run Windows Server 2003. You will then promote the server to become a domain controller in the *contoso.com* domain.

Exercise 1: Installing Windows Server 2003

This exercise should be performed on a computer compatible with Windows Server 2003. It assumes that the primary hard drive is completely empty. If your disk already has partitions configured, you can modify the exercise to match the configuration of your system.

1. Configure the computer's BIOS or the disk controller BIOS to boot from CD-ROM. If you are not sure how to configure your computer or disk controller to boot from CD-ROM, consult your hardware documentation.
2. Insert the Windows Server 2003 installation CD-ROM into the CD-ROM drive and restart the computer.
3. If the primary disk is not empty, a message appears prompting you to press any key to boot from CD. If you see this message, press any key.

   After the computer starts, a brief message appears explaining that your system configuration is being inspected, and then the Windows Setup screen appears.
4. If your computer requires special mass storage drivers that are not part of the Windows Server 2003 driver set, press F6 when prompted and provide the appropriate drivers.
5. The system prompts you to press F2 to perform an Automated System Recovery (ASR). Automated System Recovery is a new feature in Windows Server 2003 that replaces the Emergency Repair Disk feature of previous versions of Windows, and is described in Chapter 13. Do not press F2 at this time. Setup will continue.

   Notice that the gray status bar at the bottom of the screen indicates that the computer is being inspected and that files are loading. This is required to start a minimal version of the operating system.
6. If you are installing an evaluation version of Windows Server 2003, the Setup Notification screen appears informing you of this. Read the Setup Notification message, and then press Enter to continue.

   Setup displays the Welcome To Setup screen.

   Notice that, in addition to the initial installation of the operating system, you can use Windows Server 2003 Setup to repair a damaged Windows installation. The Recovery Console is described in Chapter 13.
7. Read the Welcome To Setup message, and then press Enter to continue. Setup displays the License Agreement screen.
8. Read the license agreement, pressing Page Down to scroll to the bottom of the screen.

9. Press F8 to accept the agreement.

   Setup displays the Windows Server 2003 Setup screen, prompting you to select an area of free space or an existing partition on which to install the operating system. This stage of setup provides a way for you to create and delete partitions on your hard disk.

   To complete the exercises in this book, you will need to configure a partition large enough to host the operating system installation (recommended minimum size is 3 GB) and unallocated space of at least 1 GB. The following steps assume your disk is at least 4 GB in size and is currently empty. You may make adjustments to accommodate your situation.

10. Press C to create a partition.

11. To create a 3 GB partition type **3072** in the Create Partition Of Size (In MB) box and press Enter.

12. Confirm that your partitioning is similar to that shown in Figure 1-2. Again, the recommendations for the hands-on exercises is a C: partition of at least 3 GB and 1 GB of unpartitioned space.

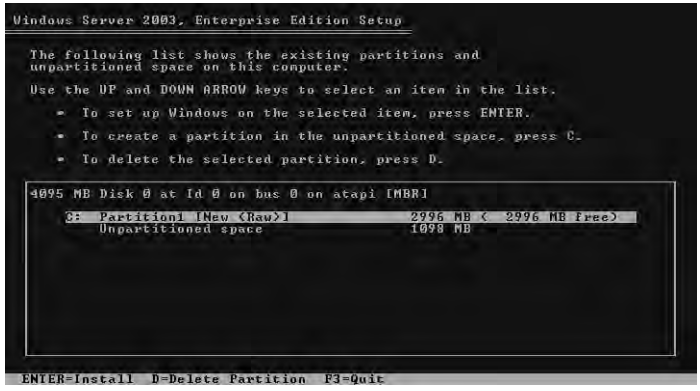

Figure 1-2 Partitioning the hard drive for setup

13. Select C: Partition1 [New (Raw)] and press Enter to install.

    You are prompted to select a file system for the partition.

14. Verify that the Format The Partition Using The NTFS File System option is selected, and press Enter to continue.

    Setup formats the partition with NTFS, examines the hard disk for physical errors that might cause the installation to fail, copies files to the hard disk, and initializes the installation. This process takes several minutes.

    Eventually, Setup displays a red status bar that counts down for 15 seconds before the computer restarts and enters the GUI mode of the setup process.

15. After the text mode of setup has completed, the system restarts. Do not, when prompted, press a key to boot to the CD-ROM.

    Windows Setup launches and produces a graphical user interface that tracks the progress of installation in the left pane. Collecting Information, Dynamic Update, and Preparing Installation options are selected. Collecting Information was completed before the GUI appeared, and Dynamic Update is not used when starting from the CD-ROM. The system is now Preparing Installation by copying files to the local disk drive.

16. On the Regional And Language Options page, choose settings that are appropriate for your language and text input requirements, and then click Next.

Tip You can modify regional settings after you install the operating system using Regional And Language Options in Control Panel.

Setup displays the Personalize Your Software page, prompting you for your name and organization name.

17. In the Name text box, type your name; in the Organization text box, type the name of an organization, and then click Next.

    Setup displays the Your Product Key page.

18. Enter the product key included with your Windows Server 2003 installation CD-ROM, and then click Next.

    Setup displays the Licensing Modes dialog box, prompting you to select a licensing mode.

19. Verify that the Per Server Number Of Concurrent Connections option is 5, and then click Next.

Important Per Server Number Of Concurrent Connections and 5 concurrent connections are suggested values to be used to complete your self-study. You should use a legal number of concurrent connections based on the actual licenses that you own. You can also choose to use Per Device Or Per User option instead of Per Server.

Setup displays the Computer Name And Administrator Password page.

Notice that Setup uses your organization name to generate a suggested name for the computer. If you didn't enter an organization name earlier in the installation process, Setup uses your name to generate part of the computer name.

20. In the Computer Name text box, type **Server01**.

    The computer name displays in all capital letters regardless of how it is entered. Throughout the rest of this self-paced training kit, the practices refer to Server01.

Caution If your computer is on a network, check with the network administrator before assigning a name to your computer.

21. In the Administrator Password text box and the Confirm Password text box, type a complex password for the Administrator account (one that others cannot easily guess). *Remember this password* because you will be logging on as Administrator to perform most hands-on exercises.

Important In a manual installation, Windows Server 2003 will not let you progress to subsequent steps until you enter an Administrator password that meets complexity requirements. You are allowed to enter a blank password, though this practice is strongly discouraged.

If the server has a modem installed, you will be presented with the Modem Dialing Information dialog box.

22. Type your area code, and then click Next. The Date And Time Settings page appears.
23. Type the correct Date & Time and Time Zone settings, and then click Next.

Important Windows Server 2003 services depend on the computer's time and date settings. Be sure to enter the correct time and date, and to select the correct time zone for your location.

24. On the Networking Settings page, select Typical Settings, and then click Next.

    The Workgroup Or Computer Domain page appears.
25. Verify that the first option is selected and that the workgroup name is Workgroup, and then click Next.

    Setup installs and configures the remaining operating system components. When the installation is complete, the computer restarts automatically and the Welcome To Windows dialog box appears.
26. Press Ctrl+Alt+Delete to initiate logon, and type the password you configured for the Administrator account.

Note Some editions of Windows Server 2003, including the Evaluation edition provided with this book, require that you activate the operating system after you install it. Activation must occur within 14 days of installation. The activation process is simple and can be completed over the Internet or by telephone. If you acquire your license to use Windows Server 2003 through one of the Microsoft volume licensing programs, you are not required to activate the license.

27. Click the balloon that appears in the System tray to initiate activation of Windows Server 2003. Follow the on-screen prompts.

Note To activate by Internet, you will have to connect Server01 to the network and you may have to adjust the TCP/IP properties of your network interface card (NIC) to reflect an appropriate IP address, subnet mask, default gateway, and DNS server address.

Exercise 2: Configuring the Server

In this exercise, you will configure the server as the first domain controller in an Active Directory domain called *contoso.com*.

Note When the Active Directory Installation Wizard is launched, the steps that it prompts you to follow will differ based on whether it detects another domain on the network. The steps presented below assume you are running the wizard on an isolated network. If you are connected to a network with another domain, the steps may vary, and you may either modify your choices appropriately or disconnect from the network prior to performing the exercise.

1. If it is not already open, open the Manage Your Server page from the Administrative Tools program group.
2. Click Add Or Remove A Role. The Configure Your Server Wizard appears.
3. Click Next and the Configure Your Server Wizard detects network settings.
4. Click Typical Configuration For A First Server, and then click Next.
5. In Active Directory Domain Name, type **contoso.com**.
6. Verify that NetBIOS Domain Name reads CONTOSO and click Next.
7. Verify that the Summary Of Selections matches that shown in Figure 1-3 and click Next.

   The Configure Your Server Wizard reminds you that the system will restart and asks you to close any open programs.
8. Click Yes.

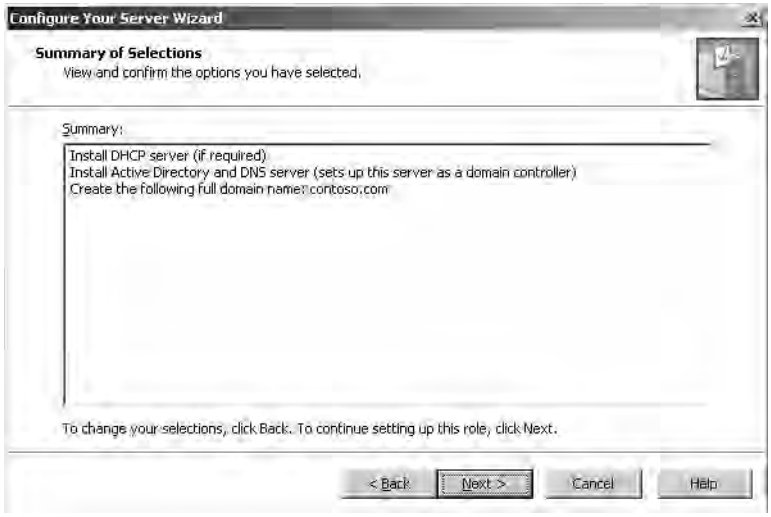

Figure 1-3 Summary Of Selections

9. After the system has restarted, log on as Administrator.

10. The Configure Your Server Wizard will summarize its final steps, as shown in Figure 1-4.

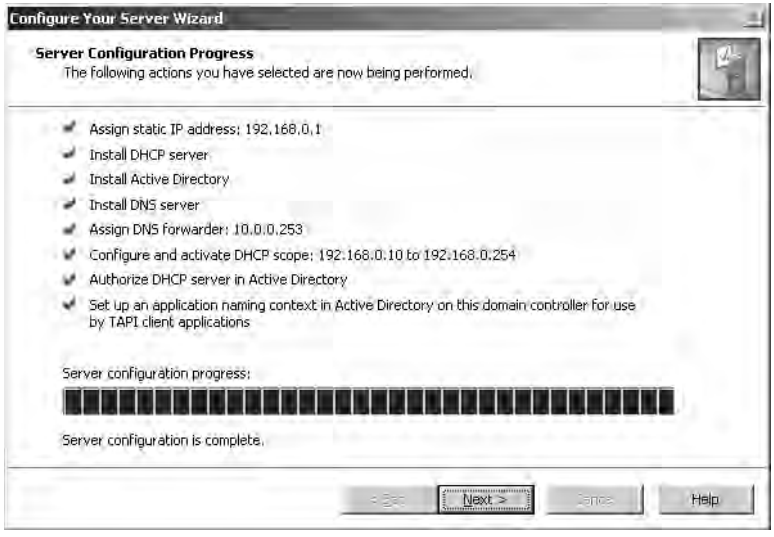

Figure 1-4 The Configure Your Server Wizard

11. Click Next and then click Finish.

12. Open Active Directory Users And Computers from the Administrative Tools group. Confirm that you now have a domain called *contoso.com* by expanding the domain and locating the computer account for Server01 in the Domain Controllers OU.

Lesson Review

1. Which of the following versions of Windows Server 2003 require product activation? (Select all that apply.)
   a. Windows Server 2003, Standard Edition, retail version
   b. Windows Server 2003, Enterprise Edition, evaluation version
   c. Windows Server 2003, Enterprise Edition, Open License version
   d. Windows Server 2003, Standard Edition, Volume License version
2. What are the distinctions among a domain, a tree, and a forest in Active Directory?

   _____

   _____

3. Which of the following is true about setup in Windows Server 2003? (Select all that apply.)
   a. Setup can be launched by booting to the CD-ROM.
   b. Setup can be launched by booting to setup floppies.
   c. Setup requires a non-blank password to meet complexity requirements.
   d. Setup will allow you to enter all 1's for the Product ID.

Lesson Summary

- Windows Server 2003 retail and evaluation versions require product activation.
- The Manage Your Server page and the Configure Your Server Wizard provide helpful guidance to the installation and configuration of additional services based on the desired server role.
- Active Directory—the Windows Server 2003 directory service—is installed on a server using the Active Directory Installation Wizard, which is launched using the Configure Your Server Wizard or by running DCPROMO from the command line.

2 Administering Microsoft Windows Server 2003

Exam Objectives in this Chapter:

- Manage servers remotely
  - Manage a server by using Remote Assistance
  - Manage a server by using Terminal Services remote administration mode
  - Manage a server by using available support tools
- Troubleshoot Terminal Services
  - Diagnose and resolve issues related to Terminal Services security
  - Diagnose and resolve issues related to client access to Terminal Services

Why This Chapter Matters

In the daily work of a systems administrator, you frequently use tools to configure user accounts, modify computer software and service settings, install new hardware, and perform many other tasks. As the computing environment expands to include more computers, so expands the amount of work to be done. The Microsoft Management Console (MMC) allows for the consolidation and organization of some of the tools used most often. In addition, MMC consoles can be customized and tailored to fit the exact needs of the worker and the task at hand, so tasks can be delegated to more junior administrators with fewer chances for error.

When more global control of a remote computer is required, beyond what can be done remotely through the MMC, two key tools make administration of remote computers possible: Remote Desktop for Administration and Remote Assistance. Generally, you can regard Remote Desktop for Administration as a client-server application that allows for a window on your desktop computer to show the local console of a server computer, giving you the ability to control the keyboard and mouse functions as if you were logged on locally at the console of the server. Remote Assistance is similar in function, but is scoped for desktop computers running an operating system from the Microsoft Windows Server 2003 or Windows XP family. A user at that computer makes a request for assistance, and a remote connection can be established from a remote computer to that desktop.

Lessons in this Chapter:

Before You Begin

To perform the practices related to the objectives in this chapter, you must have

- A computer that has Windows Server 2003 installed and operating. To follow the examples directly, your server should be named Server01 and function as a domain controller in the *contoso.com* domain.
- Remote Desktop for Administration installed on Server01, with Remote Desktop and Remote Assistance enabled.
- A configured and functioning Transmission Control Protocol/Internet Protocol (TCP/IP) network to which your console and remote administrative target computers can connect (for administration of remote computers).

Lesson 1: The Microsoft Management Console

The primary administrative tool for managing Windows Server 2003 is the MMC. The MMC provides a standardized, common interface for one or more of the applications, called *snap-ins*, that you use to configure the elements of your environment. These snap-ins are individualized to specific tasks, and can be ordered and grouped within the MMC to your administrative preference.

The primary administrative tools in Windows Server 2003 are MMC consoles with collections of snap-ins suited to a specific purpose. The Active Directory Users and Computers administrative tool, for example, is specifically designed to administer the security principals (Users, Groups, and Computers) in a domain. The snap-ins within the MMC—not the MMC itself—are the administrative tools that you use.

Note MMC consoles will run on Windows Server 2003, Windows 2000, Windows NT 4, Windows XP, and Windows 98.

After this lesson, you will be able to

- Configure an MMC with individual snap-ins
- Configure an MMC with multiple snap-ins
- Save an MMC in Author or User mode

Estimated lesson time: 15 minutes

The MMC

The MMC looks very much like a version of Windows Explorer, only with fewer buttons. The functional components of an MMC are contained within what are called snap-ins: Menus and a toolbar provide commands for manipulating the parent and child windows, and the console itself (which contains the snap-ins) allows targeted functionality. In addition, an MMC can be saved with and the various options and modes appropriate to the situation.

Navigating the MMC

An empty MMC is shown in Figure 2-1. Note that the console has a name, and that there is a Console Root. It is this Console Root that will contain any snap-ins that you choose to include.

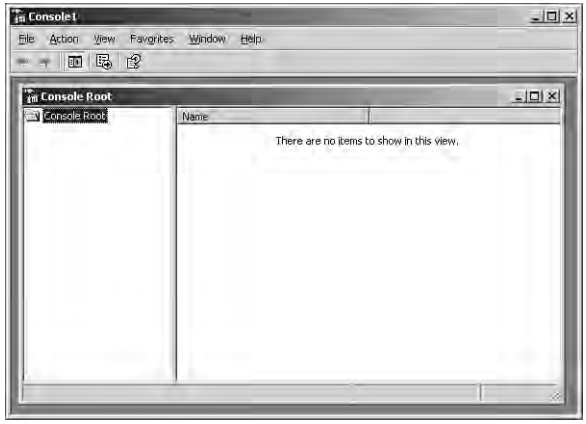

Figure 2-1 An empty MMC

Each console includes a console tree, console menu and toolbars, and the detail pane. The contents of these will vary, depending upon the design and features of the snap-in use. Figure 2-2 shows a populated MMC with two snap-ins loaded, and a child window of the Device Manager snap-in.

Figure 2-2 A populated MMC

Using the MMC Menus and Toolbar

Although each snap-in will add its unique menu and toolbar items, there are several key menus and commands that you will use in many situations that are common to most snap-ins, as shown in Table 2-1.

Table 2-1 Common MMC Menus and Commands

Menu	Commands
File	Create a new console, open an existing console, add or remove snap-ins from a console, set options for saving a console, the recent console file list, and an exit command
Action	Varies by snap-in, but generally includes export, output, configuration, and help features specific to the snap-in
View	Varies by snap-in, but includes a customize option to change general console characteristics
Favorites	Allows for adding and organizing saved consoles
Window	Open a new window, cascade, tile, and switch between open child windows in this console
Help	General help menu for the MMC as well as loaded snap-in help modules

Building a Customized MMC

Each MMC contains a collection of one or more tools called *snap-ins*. A snap-in extends the MMC by adding specific management capability and functionality. There are two types of snap-ins: stand-alone and extension.

You can combine one or more snap-ins or parts of snap-ins to create customized MMCs, which can then be used to centralize and combine administrative tasks. Although you can use many of the preconfigured consoles for administrative tasks, customized consoles allow for individualization to your needs and standardization within your environment.

Tip By creating a custom MMC, you do not have to switch between different programs or individual consoles.

Stand-Alone Snap-Ins

Stand-alone snap-ins are provided by the developer of an application. All Administrative Tools for Windows Server 2003, for example, are either single snap-in consoles or preconfigured combinations of snap-ins useful to a particular category of tasks. The Computer Management snap-in, for example, is a collection of individual snap-ins useful to a unit.

Extension Snap-Ins

Extension snap-ins, or extensions, are designed to work with one or more stand-alone snap-ins, based on the functionality of the stand-alone. When you add an extension, Windows Server 2003 places the extension into the appropriate location within the stand-alone snap-in.

Many snap-ins offer stand-alone functionality and extend the functionality of other snap-ins. For example, the Event Viewer snap-in reads the event logs of computers. If the Computer Management object exists in the console, Event Viewer automatically extends each instance of a Computer Management object and provides the event logs for the computer. Alternatively, the Event Viewer can also operate in stand-alone mode, in which case it does not appear as a node below the Computer Management node.

Off the Record Spend a few minutes analyzing your daily tasks, and group them by type of function and frequency of use. Build two or three customized consoles that contain the tools that you use most often. You will save quite a bit of time not needing to open, switch among, and close tools as often.

Console Options

Console options determine how an MMC operates in terms of what nodes in the console tree may be opened, what snap-ins may be added, and what windows may be created.

Author Mode

When you save a console in Author mode, which is the default, you enable full access to all of the MMC functionality, including:

- Adding or removing snap-ins
- Creating windows
- Creating taskpad views and tasks
- Viewing portions of the console tree
- Changing the options on the console
- Saving the console

User Modes

If you plan to distribute an MMC with specific functions, you can set the desired user mode, then save the console. By default, consoles will be saved in the Administrative Tools folder in the users' profile. Table 2-2 describes the user modes that are available for saving the MMC.

Table 2-2 MMC User Modes

Type of User Mode	Description
Full Access	Allows users to navigate between snap-ins, open windows, and access all portions of the console tree.
Limited Access, Multiple Windows	Prevents users from opening new windows or accessing a portion of the console tree, but allows them to view multiple windows in the console.
Limited Access, Single Window	Prevents users from opening new windows or accessing a portion of the console tree, and allows them to view only one window in the console.

Note MMCs, when saved, have an *.msc extension. Active Directory Users And Computers, for example, is named Dsa.msc (Directory Services Administrator.Microsoft Saved Console).

Practice: Building and Saving Consoles

In this practice you will create, configure, and save an MMC console.

Exercise 1: An Event Viewer Console

1. Click Start, and then click Run.
2. In the Open text box, type **mmc**, and then click OK.
3. Maximize the Console1 and Console Root windows.
4. From the File menu, choose Options to view the configured console mode.

   In what mode is the console running?
5. Verify that the Console Mode drop-down list box is in Author mode, and then click OK.
6. From the File menu, click Add/Remove Snap-In.

   The Add/Remove Snap-In dialog appears with the Standalone tab active. Notice that there are no snap-ins loaded.
7. In the Add/Remove Snap-In dialog box, click Add to display the Add Standalone Snap-In dialog box.
8. Locate the Event Viewer snap-in, and then click Add.

   The Select Computer dialog box appears, allowing you to specify the computer you want to administer. You can add the Event Viewer snap-in for the local computer on which you are working, or if your local computer is part of a network, you can add Event Viewer for a remote computer.

9. In the Select Computer dialog box, select Local Computer, and then click Finish.
10. In the Add Standalone Snap-In dialog box, click Close, and then in the Add/Remove Snap-Ins dialog box, click OK.

    Event Viewer (Local) now appears in the console tree. You may adjust the width of the console tree pane and expand any nodes that you want to view.
11. On your own, add a snap-in for Device Manager (local).
12. Save the MMC as MyEvents.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What is the default mode when creating an MMC?

2. Can a snap-in have focus on both the local computer and a remote computer simultaneously?

3. If you want to limit the access of a snap-in, how do you construct the MMC that contains the snap-in?

Lesson Summary

The MMC is a useful tool for organizing and consolidating snap-ins, or small programs that are used for network and computer system administrative tasks. The hierarchical display, similar to that of Windows Explorer, offers a familiar view of snap-in features in a folder-based paradigm. There are two types of snap-ins, stand-alone and extension, with extensions appearing and behaving within the MMC based on the context of their placement. Any console can be configured to work in either of two modes, Author or User, with the User mode offering some restricted functionality in the saved console.

Lesson 2: Managing Computers Remotely with the MMC

Perhaps you work in a peer-to-peer network and need to help other users create user accounts or groups on their computers to share local folders. You can save yourself a trip to your coworkers' offices by connecting to the users' computers with your Computer Management console (as shown in Figure 2-3). Or perhaps you need to format drives or perform other tasks on a remote computer. You can perform almost any task on a remote computer that you can perform locally.

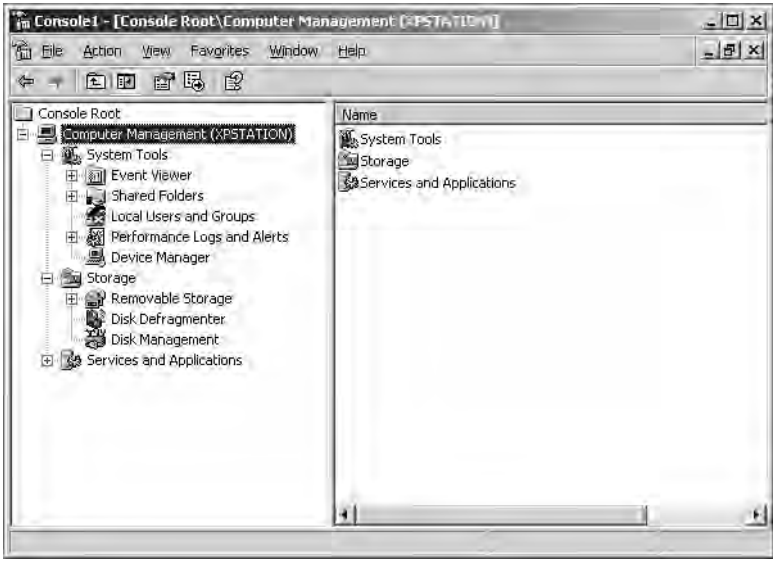

Figure 2-3 Connecting to a user's computer with the Computer Management console

After this lesson, you will be able to

- Construct an MMC to manage a computer remotely

Estimated lesson time: 10 minutes

Setting Up the Snap-In for Remote Use

To connect to and manage another system using the Computer Management console, you must launch the console with an account that has administrative credentials on the remote computer. If your credentials do not have elevated privileges on the target computer, you will be able to load the snap-in, but will not be able to read information from the target computer.

Tip You can use Run As, or secondary logon, to launch a console with credentials other than those with which you are currently logged on.

When you're ready to manage the remote system, you may open an existing console with the snap-in loaded, or configure a new MMC with a snap-in that you configure for remote connection when you build the console. If you configure an existing Computer Management console, for example, follow these steps:

1. Open the Computer Management console by right-clicking My Computer and choosing Manage from the shortcut menu.
2. Right-click Computer Management in the tree pane and choose Connect To Another Computer.
3. In the dialog box shown in Figure 2-4, type the name or IP address of the computer or browse the network for it, and then click OK to connect.

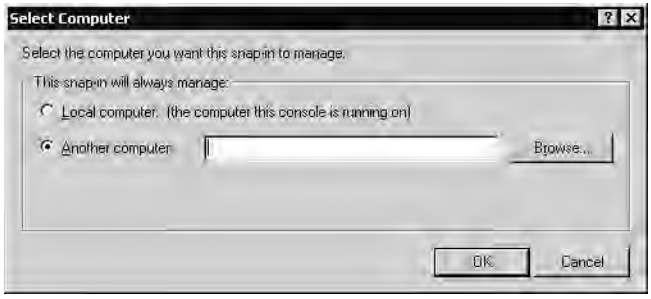

Figure 2-4 Setting the Local/Remote Context for a snap-in

Once connected, you can perform administrative tasks on the remote computer.

Practice: Adding a Remote Computer for Management (Optional)

Note This practice requires that you have a computer available for remote connection, and that you have administrative privileges on that computer.

Exercise 1: Connecting Remotely with the MMC

In this exercise, you will modify an existing MMC to connect to a remote computer.

1. Open the saved MMC from the exercise in Lesson 1 (MyEvents).
2. From the File menu, click Add/Remove Snap-In.
3. In the Add/Remove Snap-In dialog box, click Add to display the Add Standalone Snap-In dialog box.

4. Locate the Computer Management snap-in, and then click Add.
5. In the Computer Management dialog box, select Another Computer.
6. Type the name or IP address of the computer, or browse the network for it, and then click Finish to connect.
7. Click Close in the Add Standalone Snap-In dialog box, then click OK to load the Computer Management snap-in to your MyEvents console.

   You can now use the management tools to administer the remote computer.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What credentials are required for administration of a remote computer using the MMC?

2. Can an existing MMC snap-in be changed from local to remote context, or must a snap-in of the same type be loaded into the MMC for remote connection?

3. Are all functions within a snap-in used on a local computer usable when connected remotely?

Lesson Summary

The MMC is able to load many different tools in the form of snap-ins. Some of these snap-ins are programmed with the ability to connect either to the local computer or to remote computers. The connection to a remote computer can be established when the snap-in is loaded, or after loading by right-clicking the snap-in and choosing Connect. You must have administrative privileges on the remote computer to use any tools affecting the configuration of the remote computer.

Lesson 3: Managing Servers with Remote Desktop for Administration

The Windows 2000 Server family introduced a tightly integrated suite of tools and technologies that enabled Terminal Services for both remote administration and application sharing. The evolution has continued: Terminal Services is now an integral, default component of the Windows Server 2003 family, and Remote Desktop has been improved and positioned as an out-of-the-box capability, so that with one click, a Windows Server 2003 computer will allow two concurrent connections for remote administration. By adding the Terminal Server component and configuring appropriate licensing, an administrator can further extend the technologies to allow multiple users to run applications on the server. In this lesson, you will learn how to enable Remote Desktop for Administration.

After this lesson, you will be able to

- Configure a server to enable Remote Desktop for Administration
- Assign users to the appropriate group to allow them to administer servers remotely
- Connect to a server using Remote Desktop for Administration Connection

Estimated lesson time: 15 minutes

Enabling and Configuring Remote Desktop for Administration

The Terminal Services service enables Remote Desktop, Remote Assistance, and Terminal Server for application sharing. The service is installed by default on Windows Server 2003, configured in Remote Desktop for remote administration mode. Remote Desktop mode allows only two concurrent remote connections, and does not include the application sharing components of Terminal Server. Therefore, Remote Desktop operates with very little overhead on the system, and with no additional licensing requirements.

Note Because Terminal Services and its dependent Remote Desktop capability are default components of Windows Server 2003, every server has the capability to provide remote connections to its console. The term “terminal server” now therefore refers specifically to a Windows Server 2003 computer that provides application sharing to multiple users through addition of the Terminal Server component.

Other components—Terminal Server and the Terminal Server Licensing service—must be added using Add Or Remove Programs. However, all of the administrative tools required to configure and support client connections and to manage Terminal Server

are installed by default on every Windows Server 2003 computer. Each of the tools and their functions are described in Table 2-3.

Table 2-3 Default Components of Terminal Server and Remote Desktop

Installed Software	Purpose
Terminal Services Configuration	Setting properties on the Terminal Server, including session, network, client desktop, and client remote control settings
Terminal Services Manager	Sending messages to connected Terminal Server clients, disconnecting or logging off sessions, and establishing remote control or shadowing of sessions
Remote Desktop Client Installation Files	Installation of the Windows Server 2003 or Windows XP Remote Desktop Client application. The 32-bit Remote Desktop client software is installed in %*Systemroot*%\System32\Clients\Tsclient\Win32 of the Terminal Server.
Terminal Services Licensing	Configuraiton of licenses for client connections to a terminal server. This tool is not applicable for environments which utilize only Remote Desktop for Administration.

To enable Remote Desktop connections on a Windows Server 2003 computer, open the System properties from Control Panel. On the Remote tab, select Allow Users To Connect Remotely To This Computer.

Note If the Terminal Server is a Domain Controller, you must also configure the Group Policy on the Domain Controller to allow connection through Terminal Services to the Remote Desktop Users group. By default, Non-Domain Controller servers will allow Terminal Services connections by this group.

Remote Desktop Connection

Remote Desktop Connection is the client-side software used to connect to a server in the context of either Remote Desktop or Terminal Server modes. There is no functional difference from the client perspective between the two server configurations.

On Windows XP and Windows Server 2003 computers, Remote Desktop Connection is installed by default, though it is not easy to find in its default location in the All Programs\Accessories\Communications program group on the Start menu.

For other platforms, Remote Desktop Connection can be installed from the Windows Server 2003 CD or from the client installation folder (%*Systemroot*%\System32\Clients\Tsclient\Win32) on any Windows Server 2003 computer. The .msi-based Remote Desktop Connection installation package can be distributed to Windows 2000 systems using Group Policy or SMS.

Tip It is recommended to update previous versions of the Terminal Services client to the latest version of Remote Desktop Connection to provide the most efficient, secure and stable environment possible, through improvements such as a revised user interface, 128-bit encryption and alternate port selection.

Figure 2-5 shows the Remote Desktop client configured to connect to Server01 in the *contoso.com* domain.

Figure 2-5 Remote Desktop client

Configuring the Remote Desktop Client

You can control many aspects of the Remote Desktop connection from both the client and server sides. Table 2-4 lists configuration settings and their use.

Table 2-4 Remote Desktop Settings

Setting	Function
Client Settings	
General	Options for the selection of the computer to which connection should be made, the setting of static log on credentials, and the saving of settings for this connection.
Display	Controls the size of the Remote Desktop client window, color depth, and whether control-bar functions are available in full-screen mode.
Local Resources	Options to bring sound events to your local computer, in addition to standard mouse, keyboard, and screen output. How the Windows key combinations are to be interpreted by the remote computer (for example, ALT+TAB), and whether local disk, printer, and serial port connections should be available to the remote session.

Table 2-4 Remote Desktop Settings (Continued)

Setting	Function
Programs	Set the path and target folder for any program you want to start, once the connection is made.
Experience	Categories of display functions can be enabled or disabled based on available bandwith between the remote and local computers. Items include showing desktop background, showing the contents of the window while dragging, menu and window animation, themes, and whether bitmap caching should be enabled (this transmits only the changes in the screen rather than repainting the entire screen on each refresh period).
Server Settings	
Logon Settings	Static credentials can be set for the connection rather than using those provided by the client.
Sessions	Settings for ending a disconnected session, session limits and idle timeout, and reconnection allowance can be made here to override the client settings.
Environment	Overrides the settings from the user's profile for this connection for starting a program upon connection. Path and target settings set here override those set by the Remote Desktop Connection.
Permissions	Allows for additional permissions to be set on this connection.
Remote Control	Specifies whether remote control of a Remote Desktop Connection session is possible, and if it is, whether the user must grant permission at the initiation of the remote control session. Additional settings can restrict the remote control session to viewing only, or allow full interactivity with the Remote Desktop client session.
Client Settings	Override settings from the client configuration, control color depth, and disable various communication (I/O) ports.
Network Adapters	Specifies which network cards on the server will accept Remote Desktop for Administration connections.
General	Set the encryption level and authentication mechanism for connections to the server.

Terminal Services Troubleshooting

When using Remote Desktop for Administration, you are creating a connection to a server's console. There are several potential causes of failed connections or problematic sessions:

- **Network failures** Errors in standard TCP/IP networking can cause a Remote Desktop connection to fail or be interrupted. If DNS is not functioning, a client may not be able to locate the server by name. If routing is not functioning, or the Terminal Services port (by default, port 3389) misconfigured on either the client or the server, the connection will not be established.

- **Credentials** Users must belong to the Administrators or Remote Desktop Users group to successfully connect to the server using Remote Desktop for Administration.

Exam Tip Watch for group membership if access is denied when establishing a Remote Desktop for Administration connection. In earlier versions of Terminal Server, you had to be a member of the Administrators group to connect to the server, although special permissions could be established manually. Having only two remote connections to the Terminal Server is a fixed limit, and cannot be increased.

- **Policy** Domain controllers will only allow connections via Remote Desktop to administrators. You must configure the domain controller security policy to allow connections for all other remote user connections.
- **Too many concurrent connections** If sessions have been disconnected without being logged off, the server may consider its concurrent connection limit reached even though there are not two human users connected at the time. An administrator might, for example, close a remote session without logging off. If two more administrators attempt to connect to the server, only one will be allowed to connect before the limit of two concurrent connections is reached.

See Also For more on Terminal Services and the latest developments in Remote Desktop client functionality, see *http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs standard/sag_Server_Trouble_Topnode.asp*.

Practice: Installing Terminal Services and Running Remote Administration

In this practice, you will configure Server01 to enable Remote Desktop for Administration connections. You will then optimize Server01 to ensure availablity of the connection when the connection is not in use, and you will limit the number of simultaneous connections to one. You then run a remote administration session from Server02 (or another remote computer).

If you are limited to one computer for this practice, you can use the Remote Desktop client to connect to Terminal Services on the same computer. Adjust references to a remote computer in this practice to that of the local computer.

Exercise 1: Configure the Server for Remote Desktop

In this exercise, you will enable Remote Desktop connections, change the number of simultaneous connections allowed to the server, and configure the disconnection settings for the connection.

1. Logon to Server01 as Administrator.
2. Open the System properties from Control Panel.
3. On the Remote tab, enable Remote Desktop. Close System Properties.
4. Open the Terminal Services Configuration console from the Administrative Tools folder.
5. In the tscc (Terminal Services Configuration\Connections) MMC, right-click the RDP-tcp connection in the details pane, and then click Properties.
6. On the Network Adapter tab, change the Maximum Connections to 1.
7. On the Sessions tab, select both of the Override User Settings check boxes, and make setting changes so that any user session that is disconnected, by any means, or for any reason, will be closed in 15 minutes, that has no Active session time limit, and that will be disconnected after 15 minutes of inactivity.
   - ❑ End a disconnected session: 15 minutes
   - ❑ Active session limit: never
   - ❑ Idle session limit: 15 minutes
   - ❑ When session limit is reached or connection is broken: Disconnect from session

   This configuration will ensure that only one person at a time can be connected to the Terminal Server, that any disconnected session will be closed in 15 minutes, and that an idle session will be disconnected in 15 minutes. These settings are useful so as to not have a session that is disconnected or idle making the Remote Desktop for Administration connection unavailable.

Exercise 2: Connect to the Server with the Remote Desktop Client

1. On Server02 (or another remote computer, or from Server01 itself if a remote computer is not available), open Remote Desktop Connection (from the Accessories, Communications program group) and connect to and log to Server01.
2. On Server01, open the tscc (Terminal Services Configuration\Connections) MMC. You should see the remote session connected to Server01.
3. Leave the session idle for 15 minutes, or close the Remote Desktop client without logging off the Terminal Server session, and the session should be disconnected automatically in 15 minutes.

You have now logged on to Server01 remotely, and can perform any tasks on the Server01 computer that you could accomplish while logged on interactively at the console.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. How many simultaneous connections are possible to a Terminal Server running in Remote Administration mode? Why?

   _____

   _____

2. What would be the best way to give administrators the ability to administer a server remotely through Terminal Services?
   a. Don't do anything; they already have access because they are administrators.
   b. Remove the Administrators from the permission list on the Terminal Server connection, and put their administrator account in the Remote Desktop for Administration Group.
   c. Create a separate, lower-authorization user account for Administrators to use daily, and place that account in the Remote Desktop for Administration Group.
3. What tool is used to enable Remote Desktop on a server?
   a. Terminal Services Manager
   b. Terminal Services Configuration
   c. System properties in Control Panel
   d. Terminal Services Licensing

Lesson Summary

Administrators and members of the Remote Desktop Users group have the ability to connect to a server using Remote Desktop Connection. Terminal Services are installed on Windows Server 2003 by default, and allow up to two Remote Desktop for Administration connections simultaneously. The Remote Desktop Connection client, a default component of Windows XP and Windows Server 2003, can be installed on any 32-bit Windows platform from the Windows Server 2003 installation CD or (after sharing the directory) from any Windows Server 2003 computer. Configuration of Remote Desktop for Administration connections is accomplished through settings on the client (Remote Desktop Connection) and server (Terminal Server Configuration). Key settings for the connections can be overridden by the server.

Lesson 4: Using Remote Assistance

Computer users, particularly users without much technical expertise, often have configuration problems or usage questions that are difficult for a support professional or even a friend or family member to diagnose and fix over the telephone. Remote Assistance provides a way for users to get the help they need and makes it easier and less costly for corporate help desks to assist their users.

After this lesson, you will be able to

- Enable a computer to accept requests for Remote Assistance
- Use one of the available methods to request and establish a Remote Assistance session

Estimated lesson time: 30 minutes

Making the Request for Assistance

In Windows Server 2003 Help, there is a wizard-driven section for Remote Assistance, the first page of which is shown in Figure 2-6.

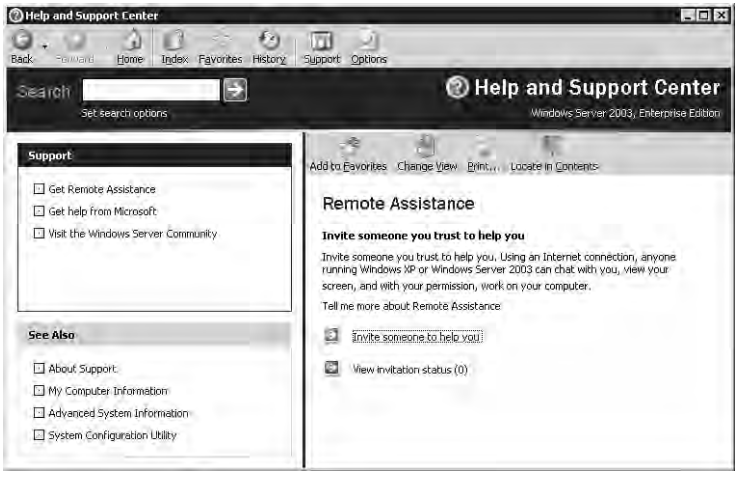

Figure 2-6 The Remote Assitance invitation screen in the Help and Support Center

The wizard-driven connection allows for a request to be sent either through a Microsoft .NET Passport account, through sending a saved file, or through a non-Passport e-mail account, along with allowing you to make a request using Windows Messenger. For a successful request through e-mail, both computers must be using a Messaging Application Programming Interface (MAPI)-compliant e-mail client.

To use the Windows Messenger service for your Remote Assistance connection, you must have the assistant's Windows Messenger user name in your contact list, and make

the request from a Windows Messenger client. Windows Messenger will display their status as online or offline. Remote Assistance can only be requested directly when your assistant is online. Remote Assistant requires that both computers are running Windows XP or a product in the Windows Server 2003 family.

Note The indicator of online status in the Remote Assistance help window is not dynamic; you must therefore refresh the screen to see an accurate status update.

After receiving a request for Remote Assistance, the helper (expert) can remotely connect to the computer and view the screen directly to fix the problem. When you initiate a request for help, the Remote Assistance client sends an encrypted ticket based on Extensible Markup Language (XML) to the helper, who is prompted to accept the invitation.

Security Alert Remote Assistance, if enabled, allows for connection to a computer under relaxed security conditions. Make certain that you provide access only to trusted authorities for Remote Assistance sessions.

Using Remote Assistance

A user can request assistance from another Windows Messenger user by placing the request through the Help and Support Center application or directly through Windows Messenger. Both applications use the same mechanisms for determining if the expert is online, and then making a request for assistance. Figure 2-7 illustrates making a request for Remote Assistance using Windows Messenger.

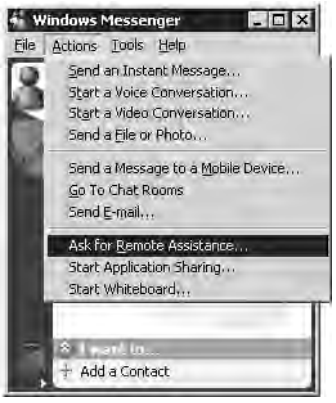

Figure 2-7 Making a request for Remote Assistance

The Windows Messenger window opens, and the user selects the expert's Windows Messenger account. The expert receives the invitation as an Instant Message. When the expert clicks Accept, the Remote Assistance session is initiated. The requesting user confirms the session by clicking Yes.

When the remote connection is established, the Remote Assistance session begins on the expert's computer. The expert and user can share desktop control, file transfer capabilities, and a chat window through which they work together to solve the user's problem.

Security Alert If the user chooses to send an e-mail or file request for Remote Assistance, a password will be required as a shared secret for the Remote Assistance session. The user should set a strong password, and let the expert know what the password is in a separate communication such as a telephone call or secure e-mail.

Offering Remote Assistance to a User

Remote Assistance is especially useful if you want to initiate troubleshooting on a user's computer. To do this, you must enable the Offer Remote Assistance Local Group Policy setting on the target (user's) local computer:

1. On the user's computer, click Start, Run, and then type **gpedit.msc**. The local Group Policy editor appears, enabling you to adjust policies that affect the local machine.

Note A Domain Group Policy may prevent you from adjusting this policy.

2. Under the Computer Configuration node, expand Administrative Templates, then System, and then click Remote Assistance.
3. Double-click Offer Remote Assistance and then select Enabled.
4. Next, click Show, then specify the individual users that will be allowed to offer assistance by assigning helpers within the context of this policy. These "helper" additions to the list should be in the form of domain\username, and must be a member of the local administrators group on the local computer.

Initializing Remote Assistance

You can now initiate Remote Assistance from your computer, to a users computer, providing that the credentials that you supply match those of a helper defined in the target computer's local Group Policy:

1. Open the Help And Support Center, click Tools, and then click Help And Support Center Tools. Next click Offer Remote Assistance. Figure 2-8 illustrates the Help And Support Center Tools interface.

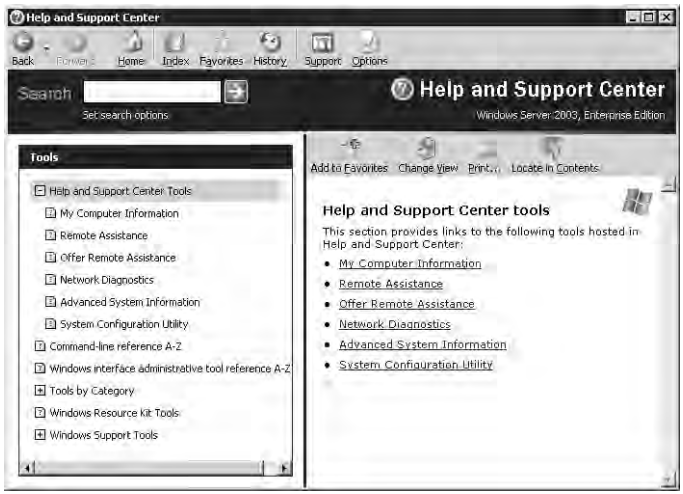

Figure 2-8 The Help And Support Center Tools

2. In the dialog box, type the name or IP address of the target computer, and then click Connect. (If prompted that several users are logged on, choose a user session.) Then click Start Remote Assistance.

   The user receives a pop-up box showing that the help-desk person is initiating a Remote Assistance session.

3. The user accepts, and Remote Assistance can proceed.

Security Alert There are several issues to consider when managing and administering Remote Assistance in the corporate environment or large organization. You can specify an open environment in which employees can receive Remote Assistance from outside the corporate firewall, or you can restrict Remote Assistance by means of Group Policy and specify various levels of permissions such as only allowing Remote Assistance from within the corporate firewall. Connections from outside the firewall require port 3389 to be open.

Firewall Constraints to Remote Assistance

Remote Assistance runs on top of Terminal Services technology, which means it must use the same port used by Terminal Services: port 3389. Remote Assistance will not work when outbound traffic from port 3389 is blocked. In addition, there are several other firewall-related concerns, particularly in relation to Network Address Translation (NAT).

- Remote Assistance supports Universal Plug and Play (UPnP) to Traverse Network Address Translation devices. This is helpful on smaller, home office networks, as Windows XP Internet Connection Sharing (ICS) supports UPnP. However, Windows 2000 ICS does *not* support UPnP.

Exam Tip Watch for questions that use Windows 2000 ICS for remote assistance from a big, corporate help desk to a small satellite office. Because Windows 2000 ICS does not support UPnP, Remote Assistance problems will abound.

- Remote Assistance will detect the Internet IP address and TCP port number on the UPnP NAT device and insert the address into the Remote Assistance encrypted ticket. The Internet IP address and TCP port number will be used to connect through the NAT device by the helper or requester workstation to establish a Remote Assistance session. The Remote Assistance connection request will then be forwarded to the client by the NAT device.
- Remote Assistance will not connect when the requester is behind a non-UPnP NAT device when e-mail is used to send the invitation file. When sending an invitation using Windows Messenger, a non-UPnP NAT device will work if one client is behind a NAT device. If both the helper and requester computers are behind non-UPnP NAT devices, the Remote Assistance connection will fail.

If you are using a software-based personal firewall or NAT in a home environment, you can use Remote Assistance with no special configurations. However, if you are using a hardware-based firewall in a home environment, the same restrictions apply: you must open port 3389 to use Remote Assistance.

Note The Instant Messenger Service itself relies upon port 1863 being open.

Practice: Using Remote Assistance through Windows Messenger

This practice requires either a partner, or a second computer for establishing the Remote Assistance session. Server01 and Server02 should have Windows Messenger installed and configured with two distinct accounts. If you are limited to a single computer for this practice, you may establish a Remote Assistance session using two separate Windows Messenger accounts configured on the same computer, but you will not be able to perform screen control.

1. From Server02 (or another computer) open Windows Messenger and log on to your Messenger Account #2.
2. From the Windows Messenger logged on as Messenger Account #1, choose Ask For Remote Assistance from the Actions menu.
3. In the Ask for Remote Assistance dialog box, choose the Messenger Account #2, and then click OK.
4. There will now be a sequence of requests and acknowledgments between the two Windows Messenger Applications. Choose Accept or OK in each query to establish the Remote Assistance session.
5. Initially, the Remote Assistance session is in Screen View Only mode. To take control of the novice's computer, you must select Take Control at the top of the Remote Assistance window. The novice user must Accept your attempt to take over the computer.

Note Either the novice or expert can end control or disconnect the session at any time.

Whether or not the expert takes over the novice's computer, screen view, file transfer, and live chat are enabled.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. How is Remote Assistance like Remote Desktop for Administration? How is it different?

_____

_____

2. What are the benefits of Remote Assistance?

3. Which of the following are firewall-related constraints relating to Remote Assistance?
   a. Port 3389 must be open
   b. NAT cannot be used
   c. Internet Connection Sharing is not possible
   d. You cannot use Remote Assistance across a Virtual Private Network (VPN)

Lesson Summary

Remote Assistance is a mutual arrangement: the user can ask an expert for help or, if properly configured through Group Policy, the expert can initiate a help session. In either case, the user must actively agree to the establishment of the session and can always give to and remove control of the user's desktop from the expert. At no time can the expert take control of the user's desktop unannounced. Remote Assistance is built upon Terminal Services and uses the interface of the help system and Windows Messenger to allow for session initiation, chat, screen viewing, screen control, and file transfer. The technology of Terminal Services and Remote Assistance is so closely tied that both services use the same network port, 3389, which must be open through any firewall for the Remote Assistance session to succeed.

Case Scenario Exercise

As part of the Remote Administration of your enterprise, your company has enabled Remote Assistance on each computer. Your sales representatives travel frequently, and use laptops to perform their work while they travel.

On your internal network, you use Windows Messenger for spontaneous communication with your clients, and for Remote Assistance. You do not, however, allow for Instant Messenger traffic across the Internet by closing port 1863 at the firewall.

You want to perform Remote Assistance for your remote users, but cannot connect to them with Windows Messenger to determine whether they are online.

Is Remote Assistance possible for your remote users? If so, how would you accomplish it?

You must use one of the alternate methods of requesting Remote Assistance.

- **The E-Mail Method** Send an e-mail to the expert through Help and Support Tools. When the expert accesses the link in the e-mail, the expert will be able to establish a Remote Assistance session.
- **File Method** Create a Remote Assistance file through Help and Support Tools. E-mail the file to the expert, or have the expert access it through a file share point. When the expert accesses the link within the file, the expert will be able to establish a Remote Assistance session.

In both methods, it is highly recommended that you create a password for the Remote Assistance session, and give the expert the password in a secure fashion so that your Remote Assistance session cannot be accessed by an unauthorized person.

Troubleshooting Lab

You are trying to connect to a Windows Server 2003 server in your environment with a Remote Desktop Connection, but consistently get the message shown in Figure 2-9 when attempting to connect.

Figure 2-9 Error Logon Message when connecting to the Remote Desktop For Administration console

You have checked settings on the server, and confirmed the following:

- You are a member of the Remote Desktop Users group.
- You are not a member of the Administrators group.
- You are able to connect to share points on the Terminal Server computer, and the computer responds affirmatively to a ping.

What other settings will you check on the Terminal Server computer to troubleshoot this problem?

It is likely that the Terminal Server in question is a Domain Controller, and that the Default Domain Controller Group Policy has not been enabled to allow remote connections by the Remote Administrative Users group. The Local Group Policy on Domain Controllers forbids non-administrator remote connections, and must be changed. The easiest way to change the Local Policy is to override it with a change to the Default Domain Controller Group Policy.

Chapter Summary

- MMCs are the common, system tool interface in Windows Server 2003.
- Snap-ins are individual tools that can be loaded into an MMC.
- Some snap-ins can be used to configure remote computers; others are limited to local computer access.
- MMCs can be saved in either Author (full access) or User (limited access) modes. The mode of an MMC does not empower or disable a user from being able to do that which they have authorization and access to do via permission sets.
- Remote Desktop for Administration allows for the same administration of a server from a remote location as if logged on to the local console interactively.
- Remote Desktop for Administration, for desktop operating systems, is available only with Windows XP.
- Remote Assistance is like Remote Desktop for Administration for the desktop, allowing remote viewing and control of Windows XP desktop computers.
- Remote Assistance will also work on a Windows Server 2003 server.
- Two users are required for Remote Assistance to be viable: one user at the target desktop, and the expert helper at another computer. Both must agree on the control actions taken during the session, and the session can be ended by either party at any time.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- MMCs are the containers for snap-ins.
- Snap-ins can be used either in local or remote context, but cannot be connected to both the local and remote computers simultaneously.
- Snap-ins can be combined in a single console to suit administrative preference.
- MMCs can be saved in User mode to restrict their configuration, but the ability to perform tasks with the tool is governed by permissions, not by limitations placed on a particular MMC console. If a user has sufficient privilege to administer a computer, the user can create MMCs with any snap-in.

- Remote Desktop for Administration requires permissions to attach with the Remote Desktop client. By default, this permission is only granted to Administrators.
- Remote Assistance is a two-way, agreed session. At no time can an expert take unauthorized control of a user's computer.
- Port 3389, the same port used by Remote Desktop for Administration, must be open at the firewall for Remote Assistance sessions to be established.

Key Terms

Remote Assistance vs. Remote Desktop for Administration Remote Assistance allows a remote control session to be established from an expert user as invited by a novice user. The credentials for authentication are supplied in the form of a shared secret password created within the invitation by the novice. Remote Desktop for Administration involves only one user connected remotely to a computer running the Terminal Server service and configured to allow Remote Desktop connections by the user.

Microsoft Management Console (MMC) What functionality is possible through remote connection of a snap-in, and what credentials are required.

Remote Desktop for Administration Credentials and server configuration required for Remote Desktop for Administration connections.

3 User Accounts

Exam Objectives in this Chapter:

- Create and manage user accounts
- Create and modify user accounts by using the Active Directory Users And Computers Microsoft Management Console (MMC) snap-in
- Create and modify user accounts by using automation
- Import user accounts
- Manage local, roaming, and mandatory user profiles
- Troubleshoot user accounts
- Diagnose and resolve account lockouts
- Diagnose and resolve issues related to user account properties
- Troubleshoot user authentication issues.

Why This Chapter Matters

Before individuals in your enterprise can access the resources they require, you must enable authentication of those individuals. Of course, the primary component of that authentication is the user's identity, maintained as an account in the Microsoft Active Directory directory service. In this chapter, you will review and enhance your knowledge related to the creation, maintenance, and troubleshooting of user accounts and authentication.

Each enterprise, and each day, brings with it a unique set of challenges related to user management. The properties you configure for a standard user account are likely to be different from those you apply to the account of a help desk team member, which are different still from those configured on the built-in Administrator account. Skills that are effective to create or modify a single user account become clumsy and inefficient when you are working with masses of accounts, such as when managing the accounts for new hires.

To address a diverse sampling of account management scenarios effectively, we will examine a variety of user management skills and tools including the Active Directory Users And Computers snap-in and powerful command-line utilities.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to user accounts in Active Directory. This training kit presumes you have a minimum of 18 months' experience and a working knowledge of Active Directory, the MMC, and the Active Directory Users And Computers snap-in. If you desire hands-on practice, using the examples and lab exercises in the chapter, prepare the following:

- A Microsoft Windows Server 2003 (Standard or Enterprise) computer installed as Server01 and configured as a domain controller in the domain *contoso.com*
- First-level organizational units (OUs): Administrative Groups, Employees, and Security Groups
- Global groups, in the Security Groups OU, called Sales Representatives and Sales Managers
- The Active Directory Users And Computers MMC, or a customized console with the Active Directory Users And Computers snap-in

Lesson 1: Creating and Managing User Objects

Active Directory requires the verification of an individual's identity—a process called authentication—before that individual can access resources. The cornerstone of authentication is the user account, with its user logon name, password, and unique security identifier (SID). During logon, Active Directory authenticates the user name and password entered by the user. The security subsystem can then build the security access token that represents that user. The access token contains the user account's SID, as well as the SIDs of groups to which the user belongs. That token can then be used to verify user rights assignments, including the right to log on locally to the system, and to authorize access to resources secured by access control lists (ACLs).

The user account is integrated into the Active Directory user object. The user object includes not just the user's name, password, and SID, but also contact information, such as telephone numbers and addresses; organizational information including job title, direct reports and manager; group memberships; and configuration such as roaming profile, terminal services, remote access, and remote control settings. This lesson will review and enhance your understanding of user objects in Active Directory.

After this lesson, you will be able to

- Create user objects in Active Directory using the Active Directory Users and Computers snap-in
- Configure user object properties
- Understand important account options that are not self-explanatory based on their descriptions
- Modify properties of multiple users simultaneously

Estimated lesson time: 15 minutes

Creating User Objects with Active Directory Users and Computers

You can create a user object with the Active Directory Users and Computers snap-in. Although user objects can be created in the domain or any of the default containers, it is best to create a user in an organizational unit, so that administrative delegation and Group Policy Objects (GPOs) can be fully leveraged.

To create a user object, select the container in which you want to create the object, click the Action menu, then choose New and choose User. You must be a member of the Enterprise Admins, Domain Admins, or Account Operators groups, or you must have been delegated administrative permissions to create user objects in the container. If you do not have sufficient permissions to create user objects, the New User command will be unavailable to you.

The New Object–User dialog box appears, as shown in Figure 3-1. The first page of the New Object–User dialog box requests properties related to the user name. Table 3-1 describes the properties that appear on the first page of the dialog box.

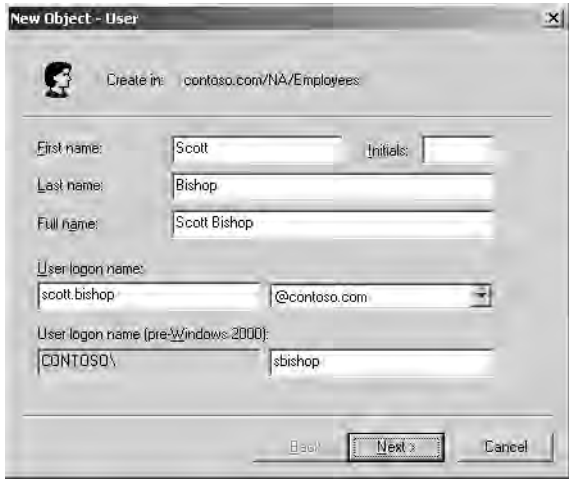

Figure 3-1 The New Object–User dialog box

Table 3-1 User Properties in the First Page of the New Object–User Dialog Box

Property	Description
First Name	The user's first name. Not required.
Initials	The middle initials of the user's name. Not required.
Last Name	The user's last name. Not required.
Full Name	The user's full name. If you enter values for the first or last name, the full name property is populated automatically. However, you can easily modify the suggested value. The field is required. The name entered here generates several user object properties, specifically CN (common name), DN (distinguished name), name, and displayName. Because CN must be unique within a container, the name entered here must be unique relative to all other objects in the OU (or other container) in which you create the user object.
User Logon Name	The user principal name (UPN) consists of a logon name and a UPN suffix which is, by default, the DNS name of the domain in which you create the object. The property is required and the entire UPN, in the format *logon-name@UPN-suffix*, must be unique within the Active Directory forest. A sample UPN would be *someone@contoso.com*. The UPN can be used to log on to any Microsoft Windows system running Windows 2000, Windows XP, or Windows Server 2003.

Table 3-1 User Properties in the First Page of the New Object–User Dialog Box (Continued)

Property	Description
User Logon Name (Pre–Windows 2000)	This logon name is used to log on from down-level clients, such as Microsoft Windows 95, Windows 98, Windows Millennium Edition (Windows Me), Windows NT 4, or Windows NT 3.51. This field is required and must be unique within the domain.

Once you have entered the values in the first page of the New Object–User dialog box, click Next. The second page of the dialog box, shown in Figure 3-2, allows you to enter the user password and to set account flags.

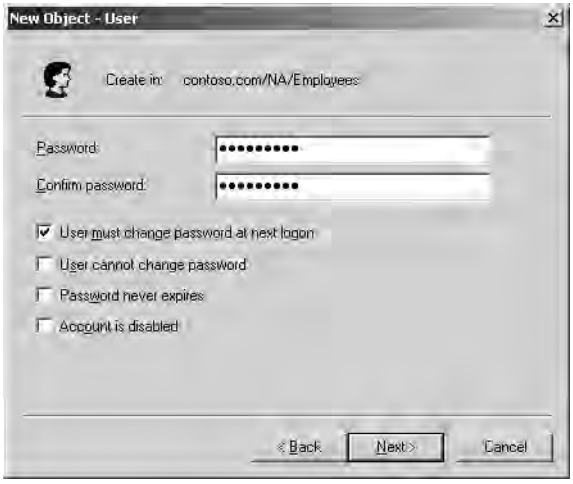

Figure 3-2 Second page of the New Object–User dialog box

Security Alert The default account policies in a Windows Server 2003 domain, set in the Default Domain Policy GPO, requires complex passwords that have a minimum of seven characters. That means a password must contain three of four character types: uppercase, lowercase, numeric, and non-alphanumeric.

When you use Windows Server 2003 in a test or lab environment, you should implement the same best practices that are required in a production network. Therefore, in this book, you are encouraged to use complex passwords for the user accounts you create; it will be left to you to remember those passwords during exercises that require logging on as those users.

The properties available in the second page of the New Object–User dialog box are summarized in Table 3-2.

Table 3-2 User Properties in the Second Page of the New Object–User Dialog Box

Property	Description
Password	The password that is used to authenticate the user. For security reasons, you should always assign a password. The password is masked as you type it.
Confirm Password	Confirm the password by typing it a second time to make sure you typed it correctly.
User Must Change Password At Next Logon	Select this check box if you want the user to change the password you have entered the first time he or she logs on. You cannot select this option if you have selected Password Never Expires. Selecting this option will automatically clear the mutually exclusive option User Cannot Change Password.
User Cannot Change Password	Select this check box if you have more than one person using the same domain user account (such as Guest) or to maintain control over user account passwords. This option is commonly used to manage service account passwords. You cannot select this option if you have selected User Must Change Password At Next Logon.
Password Never Expires	Select this check box if you never want the password to expire. This option will automatically clear the User Must Change Password At Next Logon setting, as they are mutually exclusive. This option is commonly used to manage service account passwords.
Account Is Disabled	Select this check box to disable the user account, for example, when creating an object for a newly hired employee who does not yet need access to the network.

Off the Record When creating objects for new users, choose a unique, complex password for each user that does not follow a predictable pattern. Select the option to enforce that the user must change password at next logon. If the user is not likely to log on to the network for a period, disable the account. When the user requires access to the network for the first time, ensure that the user's account is enabled. The user will be prompted to create a new, unique password that only the user knows.

Some of the account options listed in Table 3-2 have the potential to contradict policies set in the domain policies. For example, the default domain policy implements a best practice of disabling the storing of passwords using reversible encryption. However, in the rare circumstances that require reversible encryption, the user account property, Store Password Using Reversible Encryption, will take precedence for that specific user object. Similarly, the domain may specify a maximum password age, or that users must change password at next logon. If a user object is configured such that Password never expires, that configuration will override the domain's policies.

Managing User Objects with Active Directory Users And Computers

When creating a user, you are prompted to configure the most common user properties, including logon names and password. However, user objects support numerous additional properties that you can configure at any time using Active Directory Users And Computers. These properties facilitate the administration of, and the searching for, an object.

To configure the properties of a user object, select the object, click the Action menu, and then choose Properties. The user's Properties dialog box appears, as shown in Figure 3-3. An alternative way to view an object's properties would be to right-click the object and select Properties from the shortcut menu.

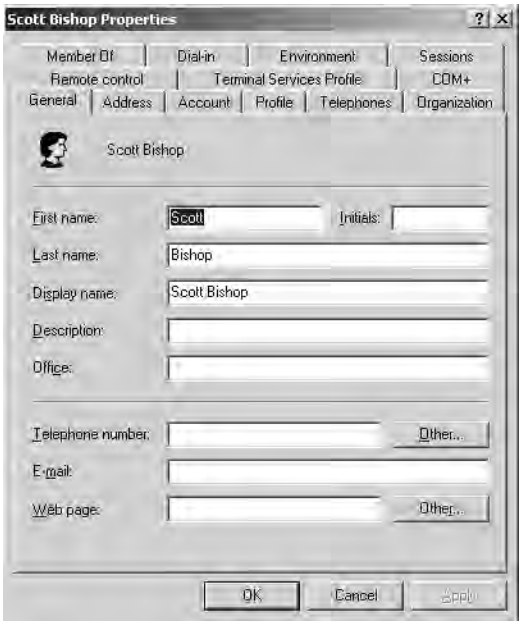

Figure 3-3 The user's Properties dialog box

The property pages in the Properties dialog box expose properties that fall into several broad categories:

- **Account properties: the Account tab** These properties include those that are configured when you create a user object, including logon names, password and account flags.
- **Personal information: the General, Address, Telephones, and Organization tabs** The General tab exposes the name properties that are configured when you create a user object.
- **User configuration management: the Profile tab** Here you can configure the user's profile path, logon script, and home folder locations.

- **Group membership: the Member Of tab** You can add and remove user groups, and set the user's primary group.
- **Terminal services: the Terminal Services Profile, Environment, Remote Control, and Sessions tabs** These four tabs allow you to configure and manage the user's experience when they are connected to a Terminal Services session.
- **Remote access: the Dial-in tab** Allows you to enable and configure remote access permission for a user.
- **Applications: the COM+ tab** Assigns Active Directory COM+ partition sets to the user. This feature, new to Windows Server 2003, facilitates the management of distributed applications.

Account Properties

Of particular note are the user's account properties, on the Account tab of the user's Properties dialog box. An example appears in Figure 3-4.

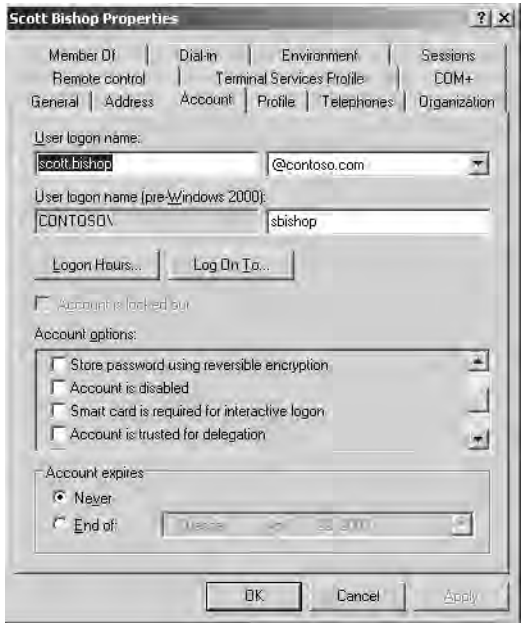

Figure 3-4 The user Account tab

Several of these properties were discussed in Table 3-2. Those properties were configured when creating the user object and can be modified, as can a larger set of account properties, using the Account tab. Several properties are not necessarily self-explanatory, and deserve definition in Table 3-3.

Table 3-3 User Account Properties

Property	Description
Logon Hours	Click Logon Hours to configure the hours during which a user is allowed to log on to the network.
Log On To	Click Log On To if you want to limit the workstations to which the user can log on. This is called Computer Restrictions in other parts of the user interface. You must have NetBIOS over TCP/IP enabled for this feature to restrict users because it uses the computer name, rather than the Media Access Control (MAC) address of its network card, to restrict logon.
Store Password Using Reversible Encryption	This option, which stores the password in Active Directory without using Active Directory's powerful, nonreversible encryption hashing algorithm, exists to support applications that require knowledge of the user password. If it is not absolutely required, do not enable this option because it weakens password security significantly. Passwords stored using reversible encryption are similar to those stored as plaintext. Macintosh clients using the AppleTalk protocol require knowledge of the user password. If a user logs on using a Macintosh client, you will need to select the option to Store password using reversible encryption.
Smart Card Is Required For Interactive Logon	Smart cards are portable, tamper-resistant hardware devices that store unique identification information for a user. They are attached to, or inserted into, a system and provide an additional, physical identification component to the authentication process.
Account Is Trusted For Delegation	This option enables a service account to impersonate a user to access network resources on behalf of a user. This option is not typically selected, certainly not for a user object representing a human being. It is used more often for service accounts in three-tier (or multi-tier) application infrastructures.
Account Expires	Use the Account Expires controls to specify when an account expires.

Managing Properties on Multiple Accounts Simultaneously

Windows Server 2003 allows you to modify the properties of multiple user accounts simultaneously. You simply select several user objects by holding the CTRL key as you click each user, or using any other multiselection options. Be certain that you select only objects of one class, such as users. Once you have multiselected, on the Action menu, choose Properties.

When you have multiselected user objects, a subset of properties is available for modification.

- **General tab** Description, Office, Telephone Number, Fax, Web Page, E-mail
- **Account tab** UPN Suffix, Logon Hours, Computer Restrictions (logon workstations), all Account Options, Account Expires
- **Address** Street, PO Box, City, State/Province, ZIP/Postal Code, Country/Region
- **Profile** Profile Path, Logon Script, and Home Folder
- **Organization** Title, Department, Company, Manager

Tip Be sure to know which properties can be modified for multiple users simultaneously. Exam scenarios that suggest a need to change many user objects' properties as quickly as possible are often testing your understanding of multiselect.

There are still many properties that must be set on a user-by-user basis. Also, certain administrative tasks, including the resetting of passwords and the renaming of accounts, can only be performed on one user object at a time.

Moving a User

If a user is transferred within an organization, it is possible that you might need to move his or her user object to reflect a change in the administration or configuration of the object. To move an object in Active Directory Users and Computers, select the object and, from the Action menu, choose Move. Alternatively, you can right-click the object and select Move from the shortcut menu.

Tip A new feature of Windows Server 2003 is that drag-and-drop operations are supported. You can move objects between OUs by dragging and dropping them in the Active Directory Users And Computers Snap-in.

Practice: Creating and Managing User Objects

In this practice, you will create three user objects. You will then modify properties of those objects.

Exercise 1: Create User Objects

1. Log on to Server01 as an administrator.
2. Open Active Directory Users And Computers.
3. Select the Employees OU.

4. Create a user account with the following information, ensuring that you use a strong password:

Text Box Name	Type
First Name	Dan
Last Name	Holme
User Logon Name	Dan.Holme
User Logon Name (Pre-Windows 2000)	Dholme

5. Create a second user object with the following properties:

Property	Type
First Name	Hank
Last Name	Carbeck
User Logon Name	Hank.Carbeck
User Logon Name (Pre-Windows 2000)	Hcarbeck

6. Create a user object for yourself, following the same conventions for user logon names as you did for the first two objects.

Exercise 2: Modify User Object Properties

1. Open the Properties dialog box for your user object.
2. Configure the appropriate properties for your user object on the General, Address, Profile, Telephones, and Organization tabs.
3. Examine the many properties associated with your user object, but do not change any other properties yet.
4. Click OK when finished.

Exercise 3: Modify Multiple User Objects' Properties

1. Open Active Directory Users And Computers and navigate to the Contoso.com Employees OU. Select the Employees OU in the tree pane, which will list the user objects you created in Exercise 1 in the details pane.
2. Click Dan Holme's user object.
3. Hold the CTRL key and click Hank Carbeck's user object.
4. Click the Action menu, and then click Properties.
5. Notice the difference between the Properties dialog box here, and the more extensive properties dialog box you explored in Exercise 2. Examine the properties that are available when multiple objects are selected, but do not modify any properties yet.
6. Configure the following properties for the two user objects:

Property Page	Property	Type
General	Description	Taught me everything I needed to know about Windows Server 2003
General	Telephone Number	(425) 555-0175
General	Web Page	*http://www.microsoft.com/mspress/*
Address	Street	One Microsoft Way
Address	City	Redmond
Address	State/Province	Washington
Address	ZIP/Postal Code	98052
Organization	Title	Author
Organization	Company	Microsoft Press

7. Click OK when you finish configuring the properties.
8. Open the properties of the object Dan Holme.
9. Confirm that the properties you configured in step 6 did, in fact, apply to the object. Click OK when you are finished.
10. Click Dan Holme's user object.
11. Hold the CTRL key and click Hank Carbeck's user object. Click the Action menu.
12. Notice that the Reset Password command is not available when you have selected more than one user object. What other commands are not available when multi-selecting? Experiment by selecting one user, opening the Action menu, then selecting two users and opening the Action menu.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You are using Active Directory Users And Computers to configure user objects in your domain, and you are able to change the address and telephone number properties of the user object representing yourself. However, the New User command is unavailable to you. What is the most likely explanation?

   _____

   _____

2. You are creating a number of user objects for a team of your organization's temporary workers. They will work daily from 9:00 A.M. to 5:00 P.M. on a contract that is scheduled to begin in one month and end two months later. They will not work outside of that schedule. Which of the following properties should you configure initially to ensure maximum security for the objects?

   a. Password

   b. Logon Hours

   c. Account expires

   d. Store password using reversible encryption

   e. Account is trusted for delegation

   f. User must change password at next logon

   g. Account is disabled

   h. Password never expires

3. Which of the following properties and administrative tasks can be configured or performed simultaneously on more than one user object?
   - **a.** Last Name
   - **b.** User Logon Name
   - **c.** Disable Account
   - **d.** Enable Account
   - **e.** Reset Password
   - **f.** Password Never Expires
   - **g.** User Must Change Password At Next Logon
   - **h.** Logon Hours
   - **i.** Computer Restrictions (Logon Workstations)
   - **j.** Title
   - **k.** Direct Reports

Lesson Summary

- You must be a member of the Enterprise Admins, Domain Admins, or Account Operators groups, or you must have been delegated administrative permissions to create user objects.
- User objects include the properties typically associated with a user "account," including logon names and password, and the unique SID for the user.
- User objects also include properties related to the individuals they represent, including personal information, group membership, and administrative settings. Windows Server 2003 allows you to change some of these properties for multiple users, simultaneously.

Lesson 2: Creating Multiple User Objects

Occasionally, situations emerge that require you to create multiple user objects quickly, such as a new class of incoming students at a school or a group of new hires at an organization. In these situations you must know how to facilitate or automate user object creation effectively so that you do not approach the task on an account-by-account basis. In Lesson 1, you learned how to create and manage user objects with Active Directory Users and Computers. This lesson will extend those concepts, skills, and tools to include user object creation through template objects, imported objects, and command-line scripting of objects.

After this lesson, you will be able to

- Create and utilize user object templates
- Import user objects from comma-delimited files
- Leverage new command-line tools to create and manage user objects

Estimated lesson time: 15 minutes

Creating and Utilizing User Object Templates

It is common for objects to share similar properties. For example, all sales representatives may belong to the same security groups, are allowed to log on to the network during the same hours, and have home folders and roaming profiles on the same server. In such cases, it is helpful when creating a user object for that object to be pre-populated with common properties. This can be accomplished by creating a generic user object—often called a *template*—and then copying that object to create new users.

To generate a user template, create a user and populate its properties. Put the user into appropriate groups.

Security Alert Be certain to *disable* the user, because it is just a template, to ensure that the account is not used for access to network resources.

To create a user based on the template, select the template and choose Copy from the Action menu. You will be prompted for properties similar to those when you created a new user: first and last name, initials, logon names, password, and account options. When the object is created, you will find that properties are copied from the template based on the following property-page-based description:

- **General** No properties copied
- **Address** All properties except Street address are copied

- **Account** All properties are copied, except for logon names, which you are prompted to enter when copying the template
- **Profile** All properties are copied, and the profile and home-folder paths are modified to reflect the new user's logon name
- **Telephones** No properties are copied
- **Organization** All properties are copied, except for Title
- **Member Of** All properties are copied
- **Dial-in, Environment, Sessions, Remote Control, Terminal Services Profile, COM+** No properties are copied

Tip A user that has been generated by copying a template has, by default, the same group membership as the template. Permissions and rights that are assigned to those groups therefore apply to the new user. However, permissions or rights assigned directly to the template user object are *not* copied or adjusted, so the new user will not have those permissions or rights.

Importing User Objects Using CSVDE

CSVDE is a command-line utility that allows you to import or export objects in Active Directory from (or to) a comma-delimited text file (also known as a comma-separated value text file), which is, of course, a common format easily read in Notepad and Microsoft Excel. The command is a powerful way to generate objects quickly. The command's basic syntax is

csvde [-i] [-f *FileName*] [-k]

-i : Specifies import mode. If not specified, the default mode is export.

-f *FileName* : Identifies the import file name.

-k : Ignores errors including "object already exists," "constraint violation," and "attribute or value already exists" during the import operation and continues processing.

The import file itself is a comma-delimited text file (*.csv or *.txt), in which the first line is a list of Lightweight Directory Access Protocol (LDAP) attribute names for the attributes imported, followed by one line for each object. Each object must contain exactly the attributes listed on the first line. A sample file follows:

```
DN,objectClass,sAMAccountName,sn,givenName,userPrincipalName

"CN=Scott Bishop,OU=Employees, DC=contoso,DC=com",
user,sbishop,Bishop,Scott,scott.bishop@contoso.com
```

This file, when imported, would create a user object in the Employees OU called Scott Bishop. The logon names, first, and last name are configured by the file. The object will be disabled initially. Once you have reset the password, you can enable the object.

See Also For more information about the powerful CSVDE command, including details regarding its parameters and its usage to *export* directory objects, open the Windows Server 2003 Help and Support Center. The LDIFDE command, also covered in detail by the Help and Support Center, allows you to import and export accounts using LDAP formats. This command and its file structure is nowhere near as intuitive for administrators as the comma-delimited file supported by CSVDE.

Utilizing Active Directory Command-Line Tools

Windows Server 2003 supports a number of powerful command-line tools to facilitate the management of Active Directory. The following is a list, and brief description, of each tool:

- **DSADD** Adds objects to the directory.
- **DSGET** Displays ("gets") properties of objects in the directory.
- **DSMOD** Modifies select attributes of an existing object in the directory.
- **DSMOVE** Moves an object from its current container to a new location.
- **DSRM** Removes an object, the complete subtree under an object, or both.
- **DSQUERY** Queries Active Directory for objects that match a specified search criteria. This command is often used to create a list of objects, which are then piped to the other command-line tools for management or modification.

These tools use one or more of the following components in their command-line switches:

- **Target object type** One of a predefined set of values that correlate with an object class in Active Directory. Common examples are: computer, user, OU, group, and server (meaning domain controller).
- **Target object identity** The distinguished name (DN) of the object against which the command is running. The DN of an object is an attribute of each object that represents the object's name and location within an Active Directory forest. For example, in Lesson 1, Exercise 1, you created a user object with the distinguished name: CN=Dan Holme, OU=Employees, DC=Contoso, DC=com.

Note When using DNs in a command parameter, enclose the name in quotes when it includes spaces. If a subcomponent of the distinguished name includes a backslash or comma, see the online help topic listed below.

- **Server** You can specify the domain controller against which you want to run the command.
- **User** You can specify a user name and password with which to run the command. This is useful if you are logged in with non-administrative credentials and wish to launch the command with elevated credentials.

In addition, switches and parameters are case-insensitive, and can be prefixed with either a dash ("-") or a slash ("/").

See Also This lesson will focus on the most commonly used commands and parameters, and on the use of these commands for user objects. For more information regarding these utilities, including the full list of parameters they accept, open the Help and Support Center and search for the phrase, "directory service command-line tools" and be sure to surround the phrase in quotes. After clicking Search, you will see the Command Line Reference on the list of Help Topics, under Search Results.

DSQUERY

The DSQUERY command queries Active Directory for objects that match a specific criteria set. The command's basic syntax is:

```
dsquery object_type  [{StartNode | forestroot | domainroot}] [-o {dn | rdn | samid}]
[-scope {subtree | onelevel | base}] [-name Name] [-desc Description] [-upn UPN]
[-samid SAMName] [-inactive NumberOfWeeks] [-stalepwd NumberOfDays] [-disabled]
[{-s Server | -d Domain}] [-u UserName] [-p {Password | *}]
```

Tip Keep in mind, this command will often be used to generate a list of objects against which you will run other command-line utilities. This is accomplished by piping the output to the second command. For example, the following command line queries Active Directory for a user object with a name starting with "Dan," pipes the result set to DSMOD, which disables each object in the result set:

dsquery user -name Dan* | dsmod user -disabled yes

The other utilities accept DNs as their input, which is the default output type as well.

The basic parameters are summarized in Table 3-4.

Table 3-4 Parameters for the DSQUERY Command

Parameter	Description
Query scope	
object_type	Required. The object type represents the object class(es) which will be searched. The object type can include computer, contact, group, OU, server, user, or the wildcard "*" to represent any object class. This lesson will focus on the command's use in querying for the user object type.
{*StartNode* forestroot | domainroot}	Optional. Specifies the node from which the search begins. You can specify the forest root (forestroot), domain root (domainroot), or a node's distinguished name (StartNode). If forestroot is specified, the search is performed using the global catalog. The default value is domainroot.
-scope {subtree | onelevel | base}	Specifies the scope of the search. A value of subtree indicates that the scope is a subtree rooted at start node. A value of onelevel indicates the immediate children of start node only. A value of base indicates the single object represented by start node. If forestroot is specified as StartNode, subtree is the only valid scope. By default, the subtree search scope is used.
How to display the result set	
-o {dn, rdn, samid}	Specifies the format in which the list of entries found by the search will be outputted or displayed. A dn value displays the distinguished name of each entry. A rdn value displays the relative distinguished name of each entry. A samid value displays the Security Accounts Manager (SAM) account name of each entry. By default, the dn format is used.
Query criteria	
-name *Name*	Searches for users whose name attributes (value of CN attribute) matches *Name*. You can use wildcards. For example, "jon*" or "*ith" or "j*th".
-desc *Description*	Searches for users whose description attribute matches *Description*. You can use wildcards.
-upn *UPN*	Searches for users whose UPN attribute matches *UPN*.
-samid *SAMName*	Searches for users whose SAM account name matches *SAMName*. You can use wildcards.
-inactive *NumberOfWeeks*	Searches for all users that have been inactive (stale) for the specified number of weeks.
-stalepwd *NumberOfDays*	Searches for all users who have not changed their passwords for the specified number of days.
-disabled	Searches for all users whose accounts are disabled.
Domain controller and credentials used for the command	
{-s *Server* | -d *Domain*}	Connects to a specified remote server or domain.

Table 3-4 Parameters for the DSQUERY Command (Continued)

Parameter	Description
-u *UserName*	Specifies the user name with which the user logs on to a remote server. By default, -u uses the user name with which the user logged on. You can use any of the following formats to specify a user name: ■ user name (for example, Linda) ■ domain\user name (for example, *widgets\Linda*) ■ UPN (for example, *Linda@widgets.microsoft.com*)
-p {*Password* | *}	Specifies to use either a password or a * to log on to a remote server. If you type *, you are prompted for a password.

Tip Inactivity is specified in weeks, but password changes are specified in days.

DSADD

The DSADD command enables you to create objects in Active Directory. When creating a user, utilize the DSADD USER command. DSADD parameters allow you to configure specific properties of an object. The parameters are self-explanatory, however the Windows Server 2003 Help And Support Center provides thorough descriptions of the DSADD command's parameters if you desire more explanation.

dsadd user *UserDN…*

The *UserDN…* parameter is one or more distinguished names for the new user object(s). If a DN includes a space, surround the entire DN with quotation marks. The *UserDN…* parameter can be entered one of the following ways:

- By piping a list of DNs from another command, such as DSQUERY.
- By typing each DN on the command line, separated by spaces.
- By leaving the DN parameter empty, at which point you can type the DNs, one at a time, at the keyboard console of the command prompt. Press ENTER after each DN. Press CTRL+Z and ENTER after the last DN.

The DSADD USER command can take the following optional parameters after the DN parameter:

- -samid *SAMName*
- -upn *UPN*
- -fn *FirstName*
- -mi *Initial*

- -ln *LastName*
- -display *DisplayName*
- -empid *EmployeeID*
- -pwd {*Password* | *} where * will prompt you for a password
- -desc *Description*
- -memberof *GroupDN*;...
- -office *Office*
- -tel *PhoneNumber*
- -email *Email*
- -hometel *HomePhoneNumber*
- -pager *PagerNumber*
- -mobile *CellPhoneNumber*
- -fax *FaxNumber*
- -iptel *IPPhoneNumber*
- -webpg *WebPage*
- -title *Title*
- -dept *Department*
- -company *Company*
- -mgr *ManagerDN*
- -hmdir *HomeDirectory*
- -hmdrv *DriveLetter*:
- -profile *ProfilePath*
- -loscr *ScriptPath*
- -mustchpwd {yes | no}
- -canchpwd {yes | no}
- -reversiblepwd {yes | no}
- -pwdneverexpires {yes | no}
- -acctexpires *NumberOfDays*
- -disabled {yes | no}

As with DSQUERY, you can add -s, -u, and -p parameters to specify the domain controller against which DSADD will run, and the user name and password—the credentials—that will be used to execute the command.

- {-s *Server* | -d *Domain*}
- -u *UserName*
- -p {*Password* | *}

The special token $username$ (case-insensitive) may replace the SAM account name in the value of the -email, -hmdir, -profile, and -webpg parameters. For example, if a SAM account name is "Denise," the -hmdir parameter can be written in either of the following formats:

- -hmdir\users\Denise\home
- -hmdir\users\$username$\home

DSMOD

The DSMOD command modifies the properties of one or more existing objects.

```
dsmod user UserDN ... parameters
```

The command handles the *UserDN…* parameter exactly as the DSADD command, and takes the same parameters. Of course now, instead of adding an object with properties, you are modifying an existing object. Note that the exceptions are that you cannot modify the *SAMName* (-samid parameter) or group membership (-memberof parameter) of a user object using the DSMOD USER command. You can use the DSMOD GROUP command, discussed in Chapter 4, "Group Accounts," to change group membership from a command-line utility.

The DSMOD command also takes the -c parameter. This parameter puts DSMOD into continuous operation mode, in which it reports errors but continues to modify the objects. Without the -c parameter, DSMOD will stop operation at the first error.

DSGET

The DSGET command gets, and outputs, selected properties of one or more existing objects.

```
dsget user UserDN ... parameters
```

The command handles the *UserDN…* parameter exactly as the DSADD command does, and takes the same parameters except that DSGET takes *only* the parameter and not an associated value. For example, DSGET takes the -samid parameter, not the -samid *SAMName* parameter and value. The reason for this is clear: You are displaying, not

adding or modifying, a property. In addition, DSGET does not support the -password parameter because it cannot display passwords. DSGET adds the -dn and -sid parameters, which display the user object's distinguished name and SID, respectively.

Exam Tip Keep track of the difference between DSQUERY and DSGET. DSQUERY finds and returns a result set of objects based on property-based search criteria. DSGET returns properties for one or more specified objects.

DSMOVE

The DSMOVE command allows you to move or rename an object within a domain. It cannot be used to move objects between domains. Its basic syntax is:

```
dsmove ObjectDN [-newname NewName] [-newparent ParentDN]
```

DSMOVE also supports the -s, -u, and -p parameters described in the section regarding DSQUERY.

The object is specified using its distinguished name in the parameter *ObjectDN*. To rename the object, specify its new common name in the *NewName* parameter. Specifying the distinguished name of a container in the *ParentDN* parameter will move the object to that container.

DSRM

DSRM is used to remove an object, its subtree, or both. The basic syntax is:

```
dsrm ObjectDN ... [-subtree [-exclude]] [-noprompt] [-c]
```

It supports the -s, -u, and -p parameters described in the section about DSQUERY.

The object is specified by its distinguished name in the *ObjectDN* parameter. The -subtree switch directs DSRM to remove the objects contents if the object is a container object. The -exclude switch excludes the object itself, and can be used only in conjunction with -subtree. Specifying -subtree and -exclude would, for example, delete an OU and its subtree, but leave the OU intact. By default, without the -subtree or -exclude switches, only the object is deleted.

You will be prompted to confirm the deletion of each object, unless you specify the -noprompt parameter. The -c switch puts DSRM into continuous operation mode, in which errors are reported but the command keeps processing additional objects. Without the -c switch, processing halts on the first error.

Practice: Creating Multiple User Objects

In this practice, you will create and manage user objects utilizing templates and command line tools.

Exercise 1: Create a User Template

1. Log on to Server01 as an administrator.
2. Open Active Directory Users And Computers.
3. Select the Employees OU in the tree pane.
4. Create a user account with the following information:

Text Box Name	Enter
First Name	Template
Last Name	Sales Representative
User Logon Name:	Template.sales.rep
User Logon Name (Pre–Windows 2000):	Templatesalesrep

5. Click Next.
6. Select Account Is Disabled. Click Next.
7. The summary page appears. Click Finish.

Note As mentioned in the chapter's "Before You Begin" section, you should create a group in the Security Groups OU called Sales Representatives. If you have not created such a group, do so now. Configure a global security group with the name Sales Representative.

8. Open the properties of the Template Sales Representative object.
9. Configure the following properties for the template account:

Tab	Property	Value
Member Of	Member Of	Sales Representatives
Account	Logon Hours	Monday–Friday, 9:00 A.M.–5:00 P.M.
Account	Expires	Three months from the current date
Organization	Company	Contoso
Profile	Profile path	\\Server1\Profiles\%Username%

10. Click OK when you have finished configuring account properties.

Exercise 2: Create Users by Copying a User Template

1. Select the Employees OU in the tree pane.
2. Select the Template Sales Representative object.
3. Click the Action menu, and then click Copy.
4. Create a new user account with the following information:

Text Box Name	Enter
First Name	Scott
Last Name	Bishop
User Logon Name:	Scott.Bishop
User Logon Name (pre-Windows 2000):	Sbishop
Account Is Disabled	Clear the check box
Password/Confirm Password	Enter and confirm a complex password as described earlier in this chapter.

5. Click Next, and then click Finish.
6. Open the properties of the object Scott Bishop.
7. Confirm that the information configured for the template on the Member Of, Account, and Organization Property pages were applied to the new object.
8. Because you will use this account for other exercises in the chapter, reset two properties. On the Account tab, set the Account Expires option to Never, and set the Logon Hours so that logon is permitted at any time.

Exercise 3: Import User Objects Using CSVDE

1. Open Notepad.
2. Type the following information carefully, creating 3 lines of text:

```
DN,objectClass,sAMAccountName,sn,givenName,userPrincipalName

"CN=Danielle Tiedt,OU=Employees,
DC=contoso,DC=com",user,dtiedt,Tiedt,Danielle,danielle.tiedt@contoso.com

"CN=Lorrin Smith-Bates,OU=Employees, DC=contoso,DC=com",user,lsmithbates,Smith-
Bates,Lorrin,lorrin.smithbates@contoso.com
```

3. Save the file as **"C:\USERS.CSV"** being certain to surround the filename with quote marks. Without quote marks, the file will be saved as C:\USERS.CSV.TXT.
4. Open the command prompt and type the following command:

   csvde –i -f c:\users.csv

5. If the command output confirms that the command completed successfully, open Active Directory Users and Computers to confirm that the objects were created. If the command output suggests that there were errors, open the USERS.CSV file in Notepad and correct the errors.
6. You will log on as these users later in this chapter. Because the users were imported without passwords, you must reset their passwords. Once the passwords have been configured, enable the accounts. Both the Reset Password and Enable Account commands can be found on either the Action or Objects shortcut menu.
7. If you have access to an application that can open comma-delimited text files, such as Microsoft Excel, open C:\USERS.CSV. You will be able to interpret its structure more easily in a columnar display than in Notepad's one-line, comma-delimited text file display.

Exercise 4: Utilize Active Directory Command-Line Tools

1. Open the command and type the following command:

   dsquery user "OU=Employees, DC=Contoso,DC=Com" -stalepwd 7

2. The command, which finds user objects that have not changed their password in seven days, should list, at a minimum, the objects you created in exercises 1 and 2. If not, create one or two new user objects and then perform step 1.
3. Type the following command and press ENTER:

   dsquery user "OU=Employees, DC=Contoso,DC=Com" -stalepwd 7 | dsmod user -mustchpwd yes

4. The command used the results of DSQUERY as the input for the DSMOD command. The DSMOD command configured the option "User must change password at next logon" for each object. Confirm your success by examining the Account tab of the affected objects.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What option will be most useful to generate 100 new user objects, each of which have identical profile path, home folder path, Title, Web Page, Company, Department, and Manager settings?

   _____

   _____

2. Which tool will allow you to identify accounts that have not been used for two months?
   a. DSADD
   b. DSGET
   c. DSMOD
   d. DSRM
   e. DSQUERY
3. What variable can be used with the DSMOD and DSADD commands to create user-specific home folders and profile folders?
   a. %*Username*%
   b. $*Username*$
   c. CN=*Username*
   d. <*Username*>
4. Which tools allow you to output the telephone numbers for all users in an OU?
   a. DSADD
   b. DSGET
   c. DSMOD
   d. DSRM
   e. DSQUERY

Lesson Summary

- A user object template is an object that is copied to produce new users. If the template is not a "real" user, it should be disabled. Only a subset of user properties are copied from templates.
- The CSVDE command enables you to import directory objects from a comma-delimited text file.
- Windows Server 2003 supports powerful new command-line tools to create, manage, and delete directory objects: DSQUERY, DSGET, DSADD, DSMOVE, DSMOD, and DSRM. Frequently, DSQUERY will produce a result set of objects that are piped as input to other commands.

Lesson 3: Managing User Profiles

You probably wouldn't read this book if you weren't supporting users, and you know that there are elements of the user's system that cause the user pain when they are not present. For example, if a user logs on and does not have access to his or her Internet Explorer Favorites, or must reconfigure his or her custom dictionary, or does not see familiar shortcuts or documents on the desktop, the user's productivity takes an instant plunge, and the help desk gets a call. Each of these examples relate to components of the user profile. Profiles can be configured to enhance their availability, security, and reliability. In this lesson, you will learn how to manage local, roaming, group, and mandatory profiles.

After this lesson, you will be able to

- Understand the application of local and roaming user profiles
- Configure a roaming user profile
- Create a preconfigured roaming user or group profile
- Configure a mandatory profile

Estimated lesson time: 15 minutes

User Profiles

A user profile is a collection of folders and data files that contain the elements of your desktop environment that make it uniquely yours. Settings include:

- Shortcuts in your Start menu, on your desktop, and in your Quick Launch bar
- Documents on your desktop and, unless redirection is configured, in your My Documents folder

Tip The properties of the My Documents folder, and the Folder Redirection policies in group policy, enable you to redirect My Documents so that it targets a network folder. This best practice allows the contents of a user's My Documents folder to be stored on a server, where they can be backed up, scanned for viruses, and made available to users throughout the organization, should they utilize a system other than their normal desktop. My Documents can also be made available offline, so that users have access to their files even when users are not connected to the network.

- Internet Explorer favorites and cookies
- Certificates (if implemented)

- Application specific files, such as the Microsoft Office custom user dictionary, user templates, and autocomplete list
- My Network Places
- Desktop display settings, such as appearance, wallpaper, and screensaver

These important elements are specific to each user. It is desirable that they are consistent between logons, available should the user need to log on to another system, and resilient in the event that the user's system fails and must be reinstalled.

Local User Profiles

By default, user profiles are stored locally on the system in the %*Systemdrive*% \Documents and Settings*%Username%* folder. They operate in the following manner:

- When a user logs on to a system for the first time, the system creates a profile for the user by copying the Default User profile. The new profile folder is named based on the logon name specified in the user's initial logon.
- All changes made to the user's desktop and software environment are stored in the local user profile. Each user has their individual profiles, so settings are user-specific.
- The user environment is extended by the All Users profile, which can include shortcuts in the desktop or start menu, network places, and even application data. Elements of the All Users profile are combined with the user's profile to create the user environment. By default, only users of the Administrators group can modify the All Users profile.
- The profile is truly local. If a user logs on to another system, the documents and settings that are part of their profile do not follow the user. Instead, the new system behaves as outlined here, generating a new local profile for the user if it is the user's first time logging on to that system.

Roaming User Profiles

If users work at more than one computer, you can configure roaming user profiles (RUPs) to ensure that their documents and settings are consistent no matter where they log on. RUPs store the profile on a server, which also means that the profiles can be backed up, scanned for viruses, and controlled centrally. Even in environments where users do not roam, RUPs provide resiliency for the important information stored in the profile. If a user's system fails and must be reinstalled, an RUP will ensure that the user's environment is identical on the new system to the one on the previous system.

To configure an RUP, create a shared folder on a server. Ideally, the server should be a file server that is frequently backed up.

Note Be sure to configure share permissions allowing Everyone Full Control. The Windows Server 2003 default share permissions allow Read, which is not sufficient for a roaming profile share.

On the Profile tab of the user's Properties dialog box, type the Profile Path in the format: ***<server>******<share>******%Username%***. The *%Username%* variable will automatically be replaced with the user's logon name.

It's that simple. The next time the user logs on, the system will identify the roaming profile location.

Exam Tip Roaming user profiles are nothing more than a shared folder and a path to the user's profile folder, within that share, entered into the user object's profile path property. Roaming profiles are not, in any way, a property of a computer object.

When the user logs *off*, the sytem will upload the profile to the profile server. The user can now log on to that system or any other system in the domain, and the documents and settings that are part of the RUP will be applied.

Note Windows Server 2003 introduces a new policy: Only Allow Local User Profiles. This policy, linked to an OU containing computer accounts, will prevent roaming profiles from being used on those computers. Instead, users will maintain local profiles.

When a user with an RUP logs on to a new system for the first time, the system does not copy its Default User profile. Instead, it downloads the RUP from the network location. When a user logs off, or when a user logs on to a system on which they've worked before, the system copies only files that have changed.

Roaming Profile Synchronization

Unlike previous versions of Microsoft Windows, Windows 2000, Windows XP, and Windows Server 2003 do not upload and download the entire user profile at logoff and logon. Instead, the user profile is *synchronized*. Only files that have changed are transferred between the local system and the network RUP folder. This means that logon and logoff with RUPs are significantly faster than with earlier Windows versions. Organizations that have not implemented RUPs for fear of their impact on logon and network traffic should reevaluate their configuration in this light.

Creating a Preconfigured User Profile

You can create a customized user profile to provide a planned, preconfigured desktop and software environment. This is helpful to achieve the following:

- Provide a productive work environment with easy access to needed network resources and applications
- Remove access to unnecessary resources and applications
- Simplify help desk troubleshooting by enforcing a more straightforward and consistent desktop

No special tools are required to create a preconfigured user profile. Simply log on to a system and modify the desktop and software settings appropriately. It's a good idea to do this as an account other than your actual user account so that you don't modify your own profile unnecessarily.

Once you've created the profile, log on to the system with administrative credentials. Open System from Control Panel, click the Advanced tab, and then click Settings in the User Profiles frame. Select the profile you created, and then click Copy To. Type the Universal Naming Convention (UNC) path to the profile in the format: **\\<server>\<share>\<username>**. In the Permitted To Use section, click Change to select the user for whom you've configured the profile. This sets the ACL on the profile folder to allow access to that user. Figure 3-5 shows an example. Click OK and the profile is copied to the network location.

Note You must be a member of the Administrators group to copy a profile.

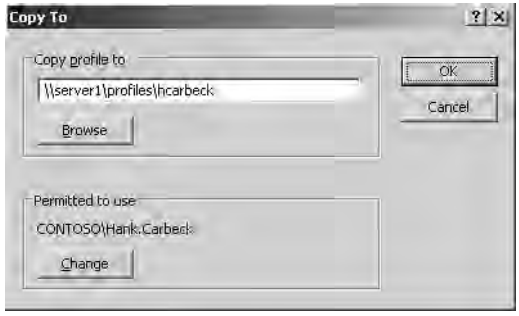

Figure 3-5 Copying a preconfigured user profile to the network

Finally, open the properties of the user object and, on the Profile tab, enter the same UNC Profile Path field. Voilà! The next time that user logs on to a domain computer, that profile will be downloaded and will determine his or her user environment.

Tip Be careful with preconfigured roaming profiles, or any roaming profiles, to pay attention to potential issues related to different hardware on systems to which a user logs on. For example, if desktop shortcuts are arranged assuming XGA (1024×768) resolution, and the user logs on to a system with a display adapter capable of only SVGA (800×600) resolution, some shortcuts may not be visible.

Profiles are also not fully cross-platform. A profile designed for Windows 98 will not function properly on a Windows Server 2003 system. You will even encounter inconsistencies when roaming between Windows Server 2003 systems and Windows XP or Windows 2000 Professional.

Creating a Preconfigured Group Profile

Roaming profiles enable you to create a standard desktop environment for multiple users with similar job responsibilities. The process is similar to creating a preconfigured user profile except that the resulting profile is made available to multiple users.

Create a profile using the steps outlined above. When copying the profile to the server, use a path such as: \\<server>\<share>\<group profile name>. You must grant access to all users who will utilize the profile, so, in the Permitted To Use frame, click Change and select a group that includes all the users, or the BUILTIN\USERS group, which includes all domain users. The only users to whom the profile will actually apply are those for which you configure the user object's profile path.

After copying the profile to the network, you must configure the profile path for the users to whom the profile will apply. Windows Server 2003 simplifies this task, in that you can multiselect users and change the profile path for all users simultaneously. Type the same UNC that you used to copy the profile to the network, for example, ***<server>****<share>****<group profile name>***.

Tip The profile path is configured as a property of one or more *user* objects. It is not assigned to a group object. Although the concept is that of a group profile, do not fall into the trap of associating the profile with a group object itself.

Finally, because more than one user will be accessing a group profile, you must make a group profile mandatory, as described in the following section.

Configuring a Mandatory Profile

A mandatory profile does not allow users to modify the profile's environment. More specifically, a mandatory profile does not maintain changes between sessions. Therefore, although a user can make changes, the next time the user logs on, the desktop will look the same as the last time he or she logged on. Changes do not persist.

Mandatory profiles can be helpful in situations in which you want to lock down the desktop. They are, in a practical sense, critical when you implement group profiles because you obviously don't want the changes one user makes to affect the environments of other users.

To configure a profile as mandatory, simply rename a file in the root folder of the profile. Interestingly, mandatory profiles are *not* configured through the application of permissions. The file you need to rename is Ntuser.dat. It is a hidden file, so you must ensure that you have specified to "Show hidden files and folders" in the Folder Options program in Control Panel, or use attrib from the command-line to remove the Hidden attribute. You may also need to configure Windows Explorer to display file extensions.

Locate the Ntuser.dat file in the profile you wish to make mandatory. Rename the file to Ntuser.man. The profile, whether roaming or local, is now mandatory.

Practice: Managing User Profiles

In this practice, you will create roaming and preconfigured roaming user profiles and mandatory group profiles. You will log on and log off a number of times. Because standard user accounts are not allowed to log on locally to a domain controller, you will begin by adding users to the Print Operators group, so that those users can log on successfully.

Exercise 1: Configure Users to Log On to the Domain Controller

In the real world, you would rarely want users to have permission to log on locally to a domain controller, however, in our one-system test environment, this capability is important. Although there are several ways to achieve this goal, the easiest is to add the Domain Users group to the Print Operators group. The Print Operators group has the right to log on locally.

1. Open Active Directory Users And Computers.
2. In the tree pane, select the Builtin container.
3. Open the Properties of the Print Operators group.
4. Use the Members tab to add Domain Users to the group.

Exercise 2: Create a Profiles Share

1. Create a Profiles folder on the C drive.
2. Right-click the Profiles folder and choose Sharing and Security.
3. Click the Sharing tab.
4. Share the folder with the default share name: Profiles.
5. Click the Permissions button.

6. Select the check box to allow Full Control.
7. Click OK.

Security Alert Windows Server 2003 applies a limited share permission by default when creating a share. Most organizations follow the best practice, which is to allow Full Control as a share permission, and to apply specific permissions to the folder using the Security tab of the folder's properties dialog box. However, in the event that an administrator has not locked down a resource before sharing it, Windows Server 2003 errs in favor of security, using a share permission that allows Read-Only access.

Exercise 3: Create a User Profile Template

1. Create a user account that will be used solely for creating profile templates. Use the following guidelines when creating the account:

Text Box Name	Enter
First Name	Profile
Last Name	Account
User Logon Name:	Profile
User Logon Name (Pre-Windows 2000):	Profile

2. Log off of Server01.
3. Log on as the Profile account.
4. Customize the desktop. You might create shortcuts to local or network resources, such as creating a shortcut to the C drive on the desktop.
5. Customize the desktop using the Display application in Control Panel. On the Desktop page of the Display Properties dialog box, you can configure the desktop background and, by clicking Customize Desktop, add the My Documents, My Computer, My Network Places, and Internet Explorer icons to the desktop.
6. Log off as the Profile account.

Exercise 4: Set Up a Preconfigured User Profile

1. Log on as Administrator.
2. Open System Properties from Control Panel, by double-clicking System.
3. Click the Advanced tab.
4. In the User Profiles frame, click Settings. This opens the Copy To dialog box.
5. Select the Profile account's user profile.

6. Click Copy To.
7. In the Copy Profile To frame, type **\\server01\profiles\hcarbeck**.
8. In the Permitted To Use section, click Change.
9. Type **Hank** and click OK.
10. Confirm the entries in the Copy To dialog box and click OK.
11. After the profile has copied to the network, click OK twice to close the User Profiles and System Properties dialog boxes.
12. Open the C:\Profiles folder to verify that the profile folder "Hcarbeck" was created.
13. Open Active Directory Users And Computers and, in the tree pane, select the Employees OU.
14. Open the properties of Hank Carbeck's user object.
15. Click the Profile tab.
16. In the Profile Path field, type **\\server01\profiles\%username%**.
17. Click Apply and confirm that the %*Username*% variable was replaced by hcarbeck. It is important that the profile path match the actual network path to the profile folder.
18. Click OK.
19. Test the success of the preconfigured roaming user profile by logging off and logging on with the user name *hank.carbeck@contoso.com*. You should see the desktop modifications that you made while logged on as the Profile account.

Exercise 5: Set Up a Preconfigured, Mandatory Group Profile

1. Log on as Administrator.
2. Open System Properties from Control Panel by double-clicking System.
3. Click the Advanced tab.
4. In the User Profiles frame, click Settings.
5. Select the Profile account's user profile.
6. Click Copy To.
7. In the Copy Profile To frame type **\\server01\profiles\sales**.
8. In the Permitted To Use frame, click Change.
9. Type **Users** and then click OK.
10. Confirm the entries in the Copy To dialog box and then click OK.

11. After the profile has copied to the network, click OK twice to close the User Profiles and System Properties dialog boxes.
12. Open the C:\Profiles folder to verify that the profile folder Sales was created.
13. Open Folder Options in Control Panel and, on the View tab, under Advanced Settings, ensure that the option, Show Hidden Files And Folders, is selected.
14. Open the C:\Profiles\Sales folder and rename the file Ntuser.dat to Ntuser.man. This makes the profile mandatory.
15. Open Active Directory Users And Computers and, in the tree pane, select the Employees OU.
16. In the details pane, select the following objects by clicking the first and pressing the CTRL key while selecting additional objects: Scott Bishop, Danielle Tiedt, Lorrin Smith-Bates.
17. Click the Action menu and choose Properties.
18. Click the Profile tab, and then select the Profile Path check box.
19. In the Profile Path field, type **\\server01\profiles\sales**.
20. Click OK.
21. Test the success of the preconfigured roaming user profile by logging off and logging on with the user name *danielle.tiedt@contoso.com*.
22. Test the mandatory nature of the profile by making a change to the desktop appearance. You will be able to make the change, but the change will not persist to future sessions.
23. Log of the computer, and then log on again as Danielle Tiedt. Because the profile is mandatory, the changes you made in the previous step should not appear.
24. Log off the computer, and log on again as Scott Bishop, with user name *scott.bishop@contoso.com*. The same desktop should appear.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Describe how a user's desktop is created when roaming user profiles are not implemented.

_____

_____

2. Arrange, in order, the steps that reflect the creation of a preconfigured roaming user profile. Use all steps provided.
   - ❑ Customize the desktop and user environment.
   - ❑ Log on as a user with sufficient permissions to modify user account properties.
   - ❑ Copy the profile to the network.
   - ❑ Create a user account so that the profile can be created without modifying any user's current profile.
   - ❑ Log on as the profile account.
   - ❑ Enter the UNC path to the profile in a user's Profile property sheet.
   - ❑ Log on as a local or domain administrator.
3. How do you make a profile mandatory?
   a. Configure the permissions on the folder's Security property sheet to deny write permission.
   b. Configure the permissions on the folders Sharing property sheet to allow only read permission.
   c. Modify the attributes of the profile folder to specify the Read Only attribute.
   d. Rename Ntuser.dat to Ntuser.man.

Lesson Summary

- Windows Server 2003 provides individual profiles for each user who logs on to the system. Profiles are stored, by default, on the local system in %*Systemdrive*% \Documents and Settings\%*Username*%.
- Roaming profiles require only a shared folder, and the profile path configured in the user object's properties.
- Preconfigured profiles are simply profiles that are copied to the profile path before the profile path is configured in the user object.
- Group profiles must be made mandatory, by renaming Ntuser.dat to Ntuser.man, so that changes made by one user do not affect other users.

Lesson 4: Securing and Troubleshooting Authentication

Once you have configured user objects, and users are authenticating against those accounts, you expose yourself to two additional challenges: security vulnerabilities, which if unaddressed could compromise the integrity of your enterprise network; and social engineering challenges, as you work to make the network, and authentication in general, friendly and reliable for users. Unfortunately, these two dynamics are at odds with each other—the more secure a network, the less usable it becomes. In this lesson, we will address issues related to user authentication. You will learn the impact of domain account policies, including password policies and account lockout policies. You will also learn how to configure auditing for logon-related events, and to perform various authentication-related tasks on user objects.

After this lesson, you will be able to

- Identify domain account policies and their impact on password requirements and authentication
- Configure auditing for logon events
- Modify authentication-related attributes of user objects

Estimated lesson time: 15 minutes

Securing Authentication with Policy

Active Directory on Windows Server 2003 supports security policies to strengthen passwords and their use within an enterprise. Of course, you must design a password policy that is sufficiently daunting to attackers while being sufficiently convenient for users, so that they do not forget passwords (resulting in increased calls to the help desk) or, worse, write down their passwords.

A system running Windows Server 2003 as a member server maintains a policy related to its local user accounts. The local security policy can be managed using the appropriately named snap-in: Local Security Policy.

You will more often be concerned with the policy that affects domain user objects. Domain account policy is managed by the Default Domain Policy. To examine and modify this policy, open Active Directory Users and Computers. Select the domain node and choose Properties from the Action menu. Click the Group Policy tab. The GPO listed as the first, or top object link is the policy object that will drive the domain account policies. It is typically, and in best practice, the Default Domain Policy. Select that policy and click Edit. The Group Policy Object Editor console opens, focused on the Default Domain policy. Navigate to Computer Configuration, Windows Settings, Security Settings, Account Policies.

Password Policy

The domain password policies enable you to protect your network against password compromise by enforcing best-practice password management techniques. The policies are described in Table 3-5.

Table 3-5 Password Policies

Policy	Description
Enforce Password History	When this policy is enabled, Active Directory maintains a list of recently used passwords, and will not allow a user to create a password that matches a password in that history. The result is that a user, when prompted to change his or her password, cannot use the same password again, and therefore cannot circumvent the password lifetime. The policy is enabled by default, with the maximum value of 24. Many IT organizations use a value of 6 to 12.
Maximum Password Age	This policy determines when users will be forced to change their passwords. Passwords that are unchanged or infrequently changed are more vulnerable to being cracked and utilized by attackers to impersonate a valid account. The default value is 42 days. IT organizations typically enforce password changes every 30 to 90 days.
Minimum Password Age	When users are required to change their passwords—even when a password history is enforced—they can simply change their passwords several times in a row to circumvent password requirements and return to their original passwords. The Minimum Password Age policy prevents this possibility by requiring that a specified number of days must pass between password changes. Of course, a password can be reset at any time in Active Directory by an administrator or support person with sufficient permissions. But the user cannot change their password more than once during the time period specified by this setting.
Minimum Password Length	This policy specifies the minimum number of characters required in a password. The default in Windows Server 2003 is seven.
Passwords Must Meet Complexity Requirements	This policy enforces rules, or filters, on new passwords. The default password filter in Windows Server 2003 (passfilt.dll) requires that a password: ■ Is not based on the user's account name. ■ Is at least six characters long. ■ Contains characters from three of the following four character types: ❑ Uppercase alphabet characters (A…Z) ❑ Lowercase alphabet characters (a…z) ❑ Arabic numerals (0…9) ❑ Nonalphanumeric characters (for example, !$#,%) Windows Server 2003 enables this policy, by default.

Note Configuring password length and complexity requirements does not affect existing passwords. These changes will affect new accounts and changed passwords after the policy is applied.

Account Lockout Policy

Account lockout refers, in its broadest sense, to the concept that after several failed logon attempts by a single user, the system should assume that an attacker is attempting to compromise the account by discovering its password and, in defense, should lock the account so no further logons may be attempted. Domain account lockout policies determine the limitations for invalid logons, expressed in a number of invalid logons in a period of time, and the requirements for an account to become unlocked, whether by simply waiting or contacting an administrator. Table 3-6 summarizes Account Lockout policies.

Table 3-6 Account Lockout Policies

Policy	Description
Account Lockout Threshold	This policy configures the number of invalid logon attempts that will trigger account lockout. The value can be in the range of 0 to 999. A value that is too low (as few as three, for example) may cause lockouts due to normal, human error at logon. A value of 0 will result in accounts never being locked out. The lockout counter is not affected by logons to locked workstations.
Account Lockout Duration	This policy determines the period of time that must pass after a lockout before Active Directory will automatically unlock a user's account. The policy is not set by default, as it is useful only in conjunction with the Account Lockout Threshold policy. Although the policy accepts values ranging from 0 to 99999 minutes, or about 10 weeks, a low setting (5 to 15 minutes) is sufficient to reduce attacks significantly without unreasonably affecting legitimate users who are mistakenly locked out. A value of 0 will require the user to contact appropriate administrators to unlock the account manually.
Reset Account Lockout Counter After	This setting specifies the time that must pass after an invalid logon attempt before the counter resets to zero. The range is 1 to 99999 minutes, and must be less than or equal to the account lockout duration.

Cross-Platform Issues

Organizations commonly implement a mix of directory service, server, and client platforms. In environments in which Windows 95, Windows 98, Windows Me, or Windows NT 4 participate in an Active Directory domain, administrators need to be aware of several issues.

- Passwords: While Windows 2000, Windows XP Professional, and Windows Server 2003 support 127-character passwords, Windows 95, Windows 98, and Windows ME support only 14-character passwords.
- Active Directory Client: The Active Directory Client can be downloaded from Microsoft's web site and installed on Windows 95, Windows 98, Windows Me, and Windows NT 4 systems. It enables those platforms running previous editions of Windows to participate in many Active Directory features available to Windows 2000 Professional or Windows XP Professional, including the following:
  - Site-awareness: a system with the Active Directory Client will attempt to log on to a domain controller in its site, rather than to any domain controller in the enterprise.
  - Active Directory Service Interfaces (ADSI): use scripting to manage Active Directory.
  - Distributed File System (Dfs): access Dfs shared resources on servers running Windows 2000 and Windows Server 2003.
  - NT LAN Manager (NTLM) version 2 authentication: use the improved authentication features in NTLM version 2.
  - Active Directory Windows Address Book (WAB): property pages
  - Active Directory search capability integrated into the Start–Find or Start–Search commands.

The following functionalites, supported on Windows 2000 Professional and Windows XP Professional, are *not* provided by the Active Directory client on Windows 95, Windows 98, and Windows NT 4:

- Kerberos V5 authentication
- Group Policy or Change and Configuration Management support
- Service principal name (SPN), or mutual authentication.

In addition, you should be aware of the following issues in mixed environments:

- Windows 98 supports passwords of up to 14 characters long. Windows 2000, Windows XP, and Windows Server 2003 can support 127-character passwords. Be aware of this difference when configuring passwords for users who log on using Windows 98.
- Without the Active Directory client, users on systems using versions of Windows earlier than Windows 2000 can change their password only if the system has access to the domain controller performing the single master operation called primary domain controller (PDC) emulator. To determine which system is the PDC emulator in a domain, open Active Directory Users And Computers, select the domain node, choose the Operations Masters command from the Action menu, and then click the PDC tab. If the PDC emulator is unavailable (that is, if it is offline or on the distant side of a downed network connection), the user cannot change his or her password.
- As you have learned in this chapter, user objects maintain two user logon name properties. The Pre-Windows 2000 logon name, or SAM name, is equivalent to the user name in Windows 95, Windows 98, or Windows NT 4. When users log on, they enter their user name and must select the domain from the Log On To box. In other situations, the user name may be entered in the format <*DomainName*>\<*UserLogonName*>.
- Users logging on using Windows 2000 or later platforms may log on the same way, or they may log on using the more efficient UPN. The UPN takes the format <*UserLogonName*>@<*UPN Suffix*>, where the UPN suffix is, by default, the DNS domain name in which the user object resides. It is not necessary to select the domain from the Log On To box when using UPN logon. In fact, the box becomes disabled as soon as you type the "@" symbol.

Auditing Authentication

If you are concerned that attacks may be taking place to discover user passwords, or to troubleshoot authentication problems, you can configure an auditing policy that will create entries in the Security log that may prove illuminating.

Audit Policies

The following policies are located in the Computer Configuration, Windows Settings, Security Settings, Local Policies, Audit Policy node of Group Policy Object Editor (or the Local Security Policy snap-in). You can configure auditing for successful or failed events.

- **Audit Account Logon Events** This policy audits each instance of user logon that involves domain controller authentication. For domain controllers, this policy is defined in the Default Domain Controllers GPO. Note, first, that this policy will

create a Security log entry on a domain controller each time a user logs on interactively or over the network using a domain account. Second, remember that to evaluate fully the results of the auditing, you must examine the Security logs on all domain controllers, because user authentication is distributed among each domain controller in a site or domain.

- **Audit Account Management** Configures auditing of activities including the creation, deletion, or modification of user, group, or computer accounts. Password resets are also logged when account management auditing is enabled.
- **Audit Logon Events** Logon events include logon and logoff, interactively or through network connection. If you have enabled Audit Account Logon Events policy for successes on a domain controller, workstation logons will not generate logon audits. Only interactive and network logons to the domain controller itself generate logon events. Account logon events are generated on the local computer for local accounts and on the domain controller for network accounts. Logon events are generated wherever the logon occurs.

Tip Keep track of the distinction between Account Logon and Logon events. When a user logs on to their workstation using a domain account, the workstation registers a Logon event and the domain controller registers an Account Logon event. When the user connects to a network server's shared folder, the server registers a Logon event and the domain controller registers an Account Logon event.

Security Event Log

Once you have configured auditing, the security logs will begin to fill with event messages. You can view these messages by selecting Security from the Event Viewer snap-in, and then double-clicking the event.

Exam Tip Remember that Account Logon events will need to be monitored on each domain controller. Logon events must be monitored on all systems.

Administering User Authentication

When users forget their passwords, are transferred or terminated, you will have to manage their user objects appropriately. The most common administrative tasks related to user account security are unlocking an account, resetting a password, disabling, enabling, renaming, and deleting user objects.

Unlocking a User Account

The account lockout policy requires that when a user has exceeded the limit for invalid logon attempts, the account is locked and no further logons can be attempted for a specified period of time, or until an administrator has unlocked the account.

To unlock a user, select the user object and, from the Action menu, choose Properties. Click the Account tab and clear the check box: Account Is Locked Out.

Resetting User Passwords

If a user forgets his or her password, you must reset the password. You do not need to know the user's old password to do so. Simply select the user object and, from the Action menu, choose the Reset Password command. Enter the new password twice to confirm the change, and as a security best practice, select the User Must Change Password At Next Logon option.

Disabling, Enabling, Renaming, and Deleting User Objects

Personnel changes may require you to disable, enable, or rename a user object. The process for doing so is similar for each action. Select the user and, from the Action menu, choose the appropriate command, as follows:

- **Disabling And Enabling A User** When a user does not require access to the network for an extended period of time, you should disable the account. Re-enable the account when the user needs to log on once again. Note that the only one of the commands to Disable or Enable will appear on the Action menu depending on the current status of the object.
- **Deleting A User** When a user is no longer part of your organization, *and there will not soon be a replacement*, delete the user object. Remember that by deleting a user, you lose its group memberships and, by deleting the SID, its rights and permissions. If you recreate a user object with the same name, it will have a different SID, and you will have to reassign rights, permissions, and group memberships.
- **Renaming A User** You will rename a user if a user changes their name, for example through marriage, or in the event that a user is no longer part of your organization, but you are replacing that user and you want to maintain the rights, permissions, group memberships, and most of the user properties of the previous user.

Tip Be certain to understand the difference between disabling and deleting an object; and between enabling and unlocking a user.

Practice: Securing and Troubleshooting Authentication

In this practice, you will configure domain auditing policies. You will then generate logon events. Finally, you will examine and troubleshoot the results of those logons.

Exercise 1: Configure Policies

1. Open Active Directory Users And Computers.
2. Select the domain node, Contoso.com
3. From the Action menu, choose Properties.
4. On the Group Policy tab, select Default Domain Policy and then click Edit.
5. Navigate to Computer Configuration, Windows Settings, Security Settings, Account Policies, and finally Account Lockout Policy.
6. Double-click the Account Lockout Duration policy.
7. Select the Define This Policy Setting check box.
8. Type **0** for the duration, then click Apply.

   The system will prompt you that it will configure the account lockout threshold and reset counter policies. Click OK.
9. Click OK to confirm the settings, and then click OK to close the Policy dialog box.
10. Confirm that the Account Lockout Duration policy is zero, the threshold is 5, and the reset counter policy is 30 minutes.
11. Close the Group Policy Object Editor window.
12. Click OK to close the Properties dialog box for the *contoso.com* domain.
13. Select the Domain Controllers container, under the domain node.
14. From the Action menu, click Properties.
15. On the Group Policy tab, select Default Domain Controllers Policy and click Edit.
16. Navigate to Computer Configuration, Windows Settings, Security Settings, Local Policies, and finally Audit Policy.
17. Double-click the Audit Account Logon Events policy.
18. Select Define These Policy Settings, select both Success and Failure, and then click OK.
19. Double-click the Audit Logon Events policy.
20. Select Define These Policy Settings, select both Success and Failure, and then click OK.

21. Double-click the Audit Account Management policy.
22. Select Define These Policy Settings, select Success, and then click OK.
23. Close the Group Policy Object Editor window.
24. Click OK to close the Properties dialog box for the Domain Controllers Properties dialog box.

Exercise 2: Generate Logon Events

1. Log off of Server01.
2. Generate two logon failure events by attempting to log on twice with the username sbishop and an *invalid* password.
3. Log on correctly as sbishop.
4. Log off.

Exercise 3: Generate Account Management Events

1. Log on as Administrator.
2. Open Active Directory Users And Computers.
3. In the tree pane, navigate to and select the Employees OU.
4. In the details pane, select Scott Bishop's user object, and then click the Action menu.
5. Click the Reset Password command.
6. Enter and confirm a new password for Scott Bishop, and then click OK.

Exercise 4: Examine Authentication Security Event Messages

1. Open the Computer Management console from the Administrative Tools group.
2. Expand Event Viewer and select Security.
3. Make sure the Category column is wide enough that you can identify the types of events that are logged.
4. Explore the events that have been generated by recent activity. Note the failed logons, the successful logons, and the resetting of Scott Bishop's password.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You enable the password complexity policy for your domain. Describe the requirements for passwords, and when those requirements will take effect.

2. To monitor potential dictionary attacks against user passwords in your enterprise, what is the single best auditing policy to configure, and what log or logs will you evaluate?

3. A user has forgotten his or her password and attempts to log on several times with an incorrect password. Eventually, the user receives a logon message indicating that the account is either disabled or locked out. The message suggests that the user contact an administrator. What must you do?
   a. Delete the user object and recreate it.
   b. Enable the user object.
   c. Unlock the user object.
   d. Reset the password for the user object.

Lesson Summary

- The Default Domain Policy drives account policies including the password and lockout policies.
- The Default Domain Controllers Policy specifies key auditing policies for domain controllers.
- Auditing for authentication generates events in each domain controller's security logs.

Chapter Summary

- You must be a member of the Enterprise Admins, Domain Admins, or Account Operators groups, or you must have been delegated administrative permissions to create user objects.
- User objects include the properties typically associated with a user "account," including logon names and password, and the unique SID for the user. They also include a number of properties related to the individuals they represent, including personal information, group membership, and administrative settings. Windows Server 2003 allows you to change some of these properties for multiple users, simultaneously.
- A user object template is an object which is copied to produce new users. If the template is not a "real" user, it should be disabled. Only a subset of user properties are copied from templates.
- The CSVDE command enables you to import directory objects from a comma-delimited text file.
- Windows Server 2003 supports powerful new command-line tools to create, manage, and delete directory objects: DSQUERY, DSGET, DSADD, DSMOVE, DSMOD, and DSRM. Frequently, DSQUERY will produce a result set of objects that can be piped as input to other commands.
- Windows Server 2003 provides individual profiles for each user who logs on to the system. Profiles are stored, by default, on the local system in *%Systemdrive%*\Documents and Settings*%Username%*.
- Roaming profiles require only a shared folder, and the profile path configured in the user object's properties.
- Preconfigured profiles are simply profiles that are copied to the profile path before the profile path is configured in the user object.
- Group profiles must be made mandatory, by renaming Ntuser.dat to Ntuser.man, so that changes made by one user do not affect other users.
- The Default Domain Policy drives account policies including the password and lockout policies, whereas the Default Domain Controllers Policy specifies key auditing policies for domain controllers.
- Auditing for authentication generates events in each domain controller's security logs.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- The group memberships or permissions, or both, required to create user accounts.
- The options at your disposal for creating or managing multiple user accounts: user templates, importing, and command-line utilities. Understand the differences among the options, and the relative strengths and weaknesses of each option.
- The properties that can be accessed or modified, or both, when creating a user, modifying a user in Active Directory Users and Computers, copying a template, querying with DSQUERY, or adding and modifying users with DSADD and DSMOD.
- The process for configuring a roaming user profile, a preconfigured roaming user profile, or a preconfigured, mandatory group profile.
- The impact of group policy on password and account lockout settings.
- How to audit authentication events.

Key Terms

User account template You might hear this referred to by other terms, but the idea is the same. A template account is used as the basis for new accounts. It is *copied* to create a new user, and some of its properties, most notably its group memberships, are copied as well.

Disabled account versus locked account An account is disabled if it has expired, or if it has been disabled by an administrator. An account is locked out if it has been subject to invalid logons beyond the threshold specified by the account lockout policy.

Mandatory profile A user profile that does not maintain modifications between sessions. A user *can* modify a mandatory profile, but users' changes are not saved when they log off. Group profiles must be made mandatory, or a change made by one user will affect all users.

4 Group Accounts

Exam Objectives in this Chapter:

- Create and manage groups
  - Create and modify groups by using the Microsoft Active Directory Users And Computers MMC snap-in
  - Identify and modify the scope of a group
  - Manage group membership
  - Create and modify groups by using automation

Why This Chapter Matters

Users, groups, and computers are the key objects in the Active Directory directory service because they allow workers, their managers, system administrators—anyone using a computer on the network—to establish their identity on the network as a security principal. Without this identification, personnel cannot gain access to the computers, applications, and data needed to do their daily work. Although it is true that the minimal identification required is that of a user and computer, management of individual user security principals becomes needlessly complicated unless users are organized into groups. Assigning permissions to hundreds of users individually is not scalable; wise use of groups makes the process of creating and administering permissions much easier.

Microsoft Windows Server 2003 has two types of groups, each with three distinct scopes. Understanding the constructions of these groups within the correct scope ensures the best use of administrative resources when creating, assigning, and managing access to resources. The possibilities of group construction also depend on whether the domain or forest in which they are created is running in the Windows Server 2003 mixed, interim, or native domain functional level. Windows Server 2003 comes with several groups already created, or built-in. You can create as many additional groups as you need.

Lessons in this Chapter:

Before You Begin

To follow and perform the practices in this chapter, you need

- A computer designated Server01 with Windows Server 2003 installed.
- Server01 should be a domain controller in the *contoso.com* domain.

Lesson 1: Understanding Group Types and Scopes

Groups are containers that can contain user and computer objects within them as members. When security permissions are set for a group in the access control list (ACL) on a resource, all members of that group receive those permissions.

Windows Server 2003 has two group types: security and distribution. *Security groups* are used to assign permissions for access to network resources. *Distribution groups* are used to combine users for e-mail distribution lists. Security groups can be used as a distribution group, but distribution groups cannot be used as security groups. Proper planning of group structure affects maintenance and scalability, especially in the enterprise environment, in which multiple domains are involved.

Tip Although settings for individual security principals—users and computers—can be set by ACLs, those settings are the exception rather than the rule of best administrative practices. If you find that you are setting an inordinate number of exceptions in ACLs for a user within a group, the user's membership in that group should be reexamined.

After this lesson, you will be able to

- Identify the two types of groups and their proper use
- Identify the three types of group scope and their proper use
- Understand the difference between groups and identities

Estimated lesson time: 15 minutes

Domain Functional Levels

In Windows Server 2003, four domain functional levels are available: Windows 2000 mixed (default), Windows 2000 native, Windows Server 2003 interim, and Windows Server 2003.

- **Windows 2000 mixed** For supporting Windows NT 4, Windows 2000, and Windows Server 2003 domain controllers
- **Windows 2000 native** For supporting Windows 2000 and Windows Server 2003 domain controllers
- **Windows Server 2003 interim** For supporting Windows NT 4 and Windows Server 2003 domain controllers
- **Windows Server 2003** For supporting Windows Server 2003 domain controllers

Limitations on group properties discussed in this chapter and elsewhere in this book will refer to these domain functional levels.

Group Scope

Group scope defines how permissions are assigned to the group members. Windows Server 2003 groups, both security and distribution groups, are classified into one of three group scopes: domain local, global, and universal.

Note Although local groups are not considered part of the group scope of Windows Server 2003, they are included for completeness.

Local Groups

Local groups (or machine local groups) are used primarily for backward compatibility with Windows NT 4. There are local users and groups on computers running Windows Server 2003 that are configured as member servers. Domain controllers do not use local groups.

- Local groups can include members from any domain within a forest, from trusted domains in other forests, and from trusted down-level domains.
- A local group has only machinewide scope; it can grant resource permissions only on the machine on which it exists.

Domain Local Groups

Domain local groups are used primarily to assign access permissions to global groups for local domain resources. Domain local groups:

- Exist in all mixed, interim and native functional level domains and forests.
- Are available domainwide only in Windows 2000 native or Windows Server 2003 domain functional level domains. Domain local groups function as a local group on the domain controllers while the domain is in mixed functional level.
- Can include members from any domain in the forest, from trusted domains in other forests, and from trusted down-level domains.
- Have domainwide scope in Windows 2000 native and Windows Server 2003 domain functional level domains, and can be used to grant resource permission on any Windows Server 2003 computer within, but not beyond, the domain in which the group exists.

Global Groups

Global groups are used primarily to provide categorized membership in domain local groups for individual security principals or for direct permission assignment (particularly in the case of a mixed or interim domain functional level domain). Often, global groups are used to collect users or computers in the same domain and share the same job, role, or function. Global groups:

- Exist in all mixed, interim, and native functional level domains and forests
- Can only include members from within their domain
- Can be made a member of machine local or domain local group
- Can be granted permission in any domain (including trusted domains in other forests and pre–Windows 2003 domains)
- Can contain other global groups (Windows 2000 native or Windows Server 2003 domain functional level only)

Universal Groups

Universal groups are used primarily to grant access to resources in all trusted domains, but universal groups can only be used as a security principal (security group type) in a Windows 2000 native or Windows Server 2003 domain functional level domain.

- Universal groups can include members from any domain in the forest.
- In Windows 2000 native or Windows Server 2003 domain functional level, universal groups can be granted permissions in any domain, including domains in other forests with which a trust exists.

Tip Universal groups can help you represent and consolidate groups that span domains, and perform common functions across the enterprise. A useful guideline is to designate widely used groups that seldom change as universal groups.

Group Conversion

The scope of a group is determined at the time of its creation. However, in a Windows 2000 native or Windows Server 2003 domain functional level domain, domain local and global groups can be converted to universal groups if the groups are not members of other groups of the same scope. For example, a global group that is a member of another global group cannot be converted to a universal group. Table 4-1 summarizes the use of Windows Server 2003 domain groups as security principals (group type: security).

Table 4-1 Group Scope and Allowed Objects

Group Scope	Allowed Objects
Windows 2000 native or Windows Server 2003 functional level domain	
Domain Local	Computer accounts, users, global groups, and universal groups from any forest or trusted domain. Domain local groups from the same domain. Nested domain local groups in the same domain.
Global	Users, computers and global groups from same domain. Nested global (in same domain), domain local, or universal groups.
Universal	Universal groups, global groups, users and computers from any domain in the forest. Nested global, domain local, or universal groups.
Windows 2000 mixed or Windows Server 2003 interim functional level domain	
Domain Local	Computer accounts, users, global groups from any domain. Cannot be nested.
Global	Only users and computers from same domain. Cannot be nested.
Universal	Not available.

Special Identities

There are also some special groups called *special identities*, that are managed by the operating system. Special identities cannot be created or deleted; nor can their membership be modified by administrators. Special identities do not appear in the Active Directory Users And Computers snap-in or in any other computer management tool, but can be assigned permissions in an ACL. Table 4-2 details some of the special identities in Windows Server 2003.

Table 4-2 Special Identities and Their Representation

Identity	Representation
Everyone	Represents all current network users, including guests and users from other domains. Whenever a user logs on to the network, that user is automatically added to the Everyone group.
Network	Represents users currently accessing a given resource over the network (as opposed to users who access a resource by logging on locally at the computer where the resource is located). Whenever a user accesses a given resource over the network, the user is automatically added to the Network group.
Interactive	Represents all users currently logged on to a particular computer and accessing a given resource located on that computer (as opposed to users who access the resource over the network). Whenever a user accesses a given resource on the computer to which they are logged on, the user is automatically added to the Interactive group.

Table 4-2 Special Identities and Their Representation (Continued)

Identity	Representation
Anonymous Logon	The Anonymous Logon group refers to any user who is using network resources, but did not go through the authentication process.
Authenticated Users	The Authenticated Users group includes all users who are authenticated into the network by using a valid user account. When assigning permissions, you can use the Authenticated Users group in place of the Everyone group to prevent anonymous access to resources.
Creator Owner	The Creator Owner group refers to the user who created or took ownership of the resource. For example, if a user created a resource, but the Administrator took ownership of it, then the Creator Owner would be the Administrator.
Dialup	The Dialup group includes anyone who is connected to the network through a dialup connection.

Caution These groups can be assigned permissions to network resources, although caution should be used when assigning some of these groups permissions. Members of these groups are not necessarily users who have been authenticated to the domain. For instance, if you assign full permissions to a share for the Everyone group, users connecting from other domains will have access to the share.

Practice: Changing the Group Type and Scope

In this practice, you get hands-on experience creating groups and modifying their scope.

Exercise 1: Creating and Modifying a Group

In this exercise, you will change the type of group and its scope.

1. In Active Directory Users And Computers, create a global distribution group in the Users container called Agents.
2. Right-click the Agents group, and then choose Properties.

   Can you change the scope and type of the group? If not, why not?

   If you cannot change the type and scope of the group, the domain in which you are operating is still in mixed or Windows Server 2003 interim domain functional level. You must raise the domain functional level to either Windows 2000 native or Windows Server 2003 to change group type or scope.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What type of domain group is most like the local group on a member server? How are they alike?

2. If you are using universal groups in your domain or forest, and you need to give permission-based access to the members of the universal group, what configuration must be true of the universal group?

3. In a domain running in Windows Server 2003 domain functional level, what security principals can be a member of a global group?

Lesson Summary

- There are two types of groups: security and distribution. Security groups can be assigned permissions, while distribution groups are used for query containers, such as e-mail distribution groups, and cannot be assigned permissions to a resource.
- Security permissions for a group are assigned in an ACL just as any other security principal, such as a user or computer.
- In Windows 2000 native or Windows Server 2003 domain functional level, groups of both security and distribution type can be constructed as domain local, global, or universal, each with a different scope as to which security principals they can contain.

Lesson 2: Managing Group Accounts

The Active Directory Users And Computers MMC is the primary tool you will use to administer security principals—users, groups, and computers—in the domain. In the creation of groups, you will configure the scope, type, and membership for each. You will also use the Active Directory Users And Computers MMC to modify membership of existing groups.

After this lesson, you will be able to

- Create a group
- Modify the membership of a group
- Find the domain groups to which a user belongs

Estimated lesson time: 10 minutes

Creating a Security Group

The tool that you will use most often in the creation of groups is the Active Directory Users And Computers MMC, which can be found in the Administrative Tools folder. From within the Active Directory Users And Computers MMC, right-click the details pane of the container within which you want to create the group, and choose New, Group. You then must select the type and scope of group that you want to create.

The primary type of group that you will likely create is a security group because this is the type of group used to set permissions in an ACL. In a mixed or interim domain functional level domain, you can only set a security group for the domain local and global scopes. As Figure 4-1 illustrates, you cannot create a security group that has universal scope in mixed or interim domain functional level domains.

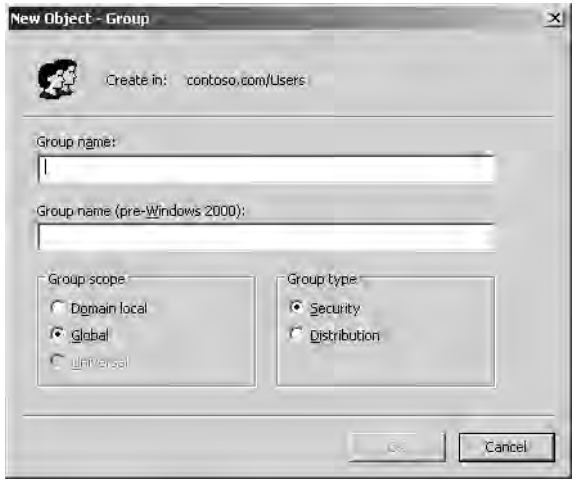

Figure 4-1 Security groups in mixed or interim functional level domains

Domain local, global, and universal groups can, however, be created as a distribution type in a mixed or interim domain functional level domain. In a mixed or interim domain functional level domain, security groups can be created in any scope.

Modifying Group Membership

Adding or deleting members from a group is also accomplished through Active Directory Users And Computers. Right-click any group, and choose Properties. Figure 4-2 illustrates the Properties dialog box of a global security group called Sales.

Figure 4-2 Properties page of the Sales security group

Table 4-3 explains the member configuration tabs of the Properties dialog box.

Table 4-3 Membership Configuration

Tab	Function
Members	Adding, removing, or listing the security principals that this container holds as members
Member Of	Adding, removing, or listing the containers that hold this container as a member

See Also See Chapter 3, "User Accounts," for additional information on using Directory Service command-line tools for viewing and modifying group membership. These tools include DSQUERY, DSGET, DSMOD, and DSGROUP. DSGET is particularly useful for listing all group memberships for a user.

Finding the Domain Groups to Which a User Belongs

Active Directory allows for flexible and creative group nesting, where

- Global groups can nest into other global groups, universal groups, or domain local groups.
- Universal groups can be members of other universal groups or domain local groups.
- Domain local groups can belong to other domain local groups.

This flexibility brings with it the potential for complexity, and without the right tools, it would be difficult to know exactly which groups a user belongs to, whether directly or indirectly. Fortunately, Windows Server 2003 adds the DSGET command, which solves the problem. From a command prompt, type:

```
dsget user UserDN -memberof [-expand]
```

The -memberof switch returns the value of the MemberOf attribute, showing the groups to which the user directly belongs. By adding the -expand switch, those groups are searched recursively, producing an exhaustive list of all groups to which the user belongs in the domain.

Practice: Modifying Group Membership

In this practice, you will work with group memberships and nesting to identify which combinations of group memberships are possible.

Exercise 1: Nesting Group Memberships

1. If the domain functional level is not already set to Windows Server 2003, use the Active Directory Users And Computers MMC to raise the domain functional level to Windows Server 2003.
2. Create three global groups in the Users Organizational Unit (OU): Group 1, Group 2, and Group 3.
3. Create three user accounts: User 1, User 2, and User 3.
4. Make User 1, User 2, and User 3 members of Group 1.
5. Make Group 1 a member of Group 2.

   Which groups can now be converted to universal groups? Test your theory (you should be able to convert 2 of the 3 groups without error).

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. In the properties of a group, which tab will you access to add users to the group?

2. You want to nest the IT Administrators group responsible for the Sales group inside the Sales group so that its members will have access to the same resources (set by permissions in an ACL) as the Sales group. From the Properties page of the IT Administrators group, what tab will you access to make this setting?

3. If your environment consists of two domains, one Windows Server 2003 and one Windows NT 4, what group scopes can you use for assigning permissions on any resource on any domain-member computer?

Lesson Summary

- Modifying group memberships is accomplished through Active Directory Users And Computers.
- If you access the properties of a security principal that is to be a member of a group, you set the group membership in the Members Of tab of the Security principal's properties. If you access the container (group) that is to hold members, set the members of the container on the Members tab.
- Groups can be nested when the domain in which they reside is set to either the Windows 2000 native or Windows Server 2003 domain functional level. If the domain is in mixed or interim domain functional level, which means that you are still supporting Windows NT 4 domain controllers, no group nesting is possible.
- Changing the type or scope of a group is only possible when the domain functional level is Windows 2000 native or Windows Server 2003.

Lesson 3: Using Automation to Manage Group Accounts

Although the Active Directory Users And Computers MMC is a convenient way to create and modify groups individually, it is not the most efficient method for creating large numbers of security principals. A tool included with Windows Server 2003, Ldifde.exe, facilitates the importing and exporting of larger numbers of security principals, including groups.

After this lesson, you will be able to

- Import security principals with LDIFDE
- Export security principles with LDIFDE
- Use the DSADD and DSMOD commands to create and modify groups

Estimated lesson time: 30 minutes

Using LDIFDE

The Lightweight Directory Access Protocol (LDAP) Data Interchange Format (LDIF) is a draft Internet standard for a file format that may be used to perform batch operations against directories that conform to the LDAP standards. LDIF can be used to export and import data, allowing batch operations such as add, create, and modify to be performed against the Active Directory. A utility program called LDIFDE is included in Windows Server 2003 to support batch operations based on the LDIF file format standard.

LDIFDE is a command-line utility, available on all Windows Server 2003 editions. From a command prompt or command shell, you run the LDIFDE utility with the appropriate command switches. Figure 4-3 lists the primary commands used with LDIFDE displayed by typing **ldifde /?** at the command prompt.

Figure 4-3 LDIFDE command-line help file

Table 4-4 details the primary LDIFDE commands.

Table 4-4 LDIFDE Commands (Primary)

Command	Usage
General parameters	
-i	Turn on Import mode (The default is Export)
-f *filename*	Input or Output *filename*
-s *servername*	The server to bind to
-c *FromDN ToDN*	Replace occurrences of FromDN to ToDN
-v	Turn on Verbose mode
-j *path*	Log File Location
-t *port*	Port Number (default = 389)
-?	Help
Export specific parameters	
-d *RootDN*	The root of the LDAP search (Default to Naming Context)
-r *Filter*	LDAP search filter (Default to "(objectClass=*)")
-p *SearchScope*	Search Scope (Base/OneLevel/Subtree)
-l *list*	List of attributes (comma-separated) to look for in an LDAP search
-o *list*	List of attributes (comma-separated) to omit from input
-g	Disable Paged Search
-m	Enable the Security Accounts Manager (SAM) logic on export
-n	Do not export binary values
Import specific parameters	
-k	The import will ignore "Constraint Violation" and "Object Already Exists" errors
Credentials parameters	
-a *UserDN*	Sets the command to run using the supplied user distinguished name and password. For example: "cn=administrator,dc=contoso,dc-com password"
-b *UserName Domain*	Sets the command to run as username domain password. The default is to run using the credentials of the currently logged on user.

Note The LDIFDE utility is included in Windows Server 2003, and can be copied to a computer running Windows 2000 Professional or Windows XP. It can then be bound and used remotely to the Windows Server 2003 Active Directory.

Real World Account Creation

Often, you will have a collection of data that already has a great deal of the information with which you will populate your Windows Server 2003 Active Directory. The data may be in a down-level domain (Windows NT 4, Windows 2000, Novell Directory Services (NDS), or some other type of database (Human Resource departments are famous for compiling data).

If you have this user data available, you can use it to populate the bulk of your Active Directory. There are many tools that are available to facilitate the extraction of data: Addusers for Windows NT 4 and LDIFDE for Windows 2000, for example. In addition, most database programs have the built-in capacity to export their data into a Comma-Separated-Value (CSV) file, which LDIFDE can import. For CSV files, however, it should be noted that some elements in object creation are mandatory, and errors will result during the import if elements are missing from the file. Group creation, however, has only the required elements of a distinguished name (CN=User) and location (DC=Domain, DC=OU), which you are unlikely to omit.

With a little editing, you can add the OU and group data to the import file, and use LDIFDE to build your Active Directory much more quickly.

Creating Groups with DSADD

The DSADD command, introduced in Chapter 2, is used to add objects to Active Directory. To add a group, use the syntax

dsadd group *GroupDN…*

The *GroupDN…* parameter is one or more distinguished names for the new user objects. If a DN includes a space, surround the entire DN with quotation marks. The *GroupDN…* parameter can be entered one of the following ways:

- By piping a list of DNs from another command, such as dsquery.
- By typing each DN on the command line, separated by spaces.
- By leaving the DN parameter empty, at which point you can type the DNs, one at a time, at the keyboard console of the command prompt. Press ENTER after each DN. Press CTRL+Z and ENTER after the last DN.

The DSADD GROUP command can take the following optional parameters after the DN parameter:

- -secgrp {*yes* | *no*} determines whether the group is a security group (yes) or a distribution group (no). The default value is yes.

- -scope {*l* | *g* | *u*} determines whether the group is a domain local (l), global (g, the default), or universal (u).
- -samid *SAMName*
- desc *Description*
- -memberof *GroupDN*... specifies groups to which to add the new group.
- -members *MemberDN*... specifies members to add to the group.

As discussed in Chapter 3, you can add -s, -u, and -p parameters to specify the domain controller against which DSADD will run, and the user name and password—the credentials—that will be used to execute the command.

- {-s *Server* | -d *Domain*}
- -u *UserName*
- -p {*Password* | *}

Modifying Groups with DSMOD

The DSMOD command, introduced in Chapter 2, is used to modify objects in Active Directory. To modify a group, use the syntax

dsmod group *GroupDN…*

The command takes many of the same switches as DSADD, including -samid, -desc, -secgrp, and -scope. Typically, though, you won't be changing those attributes of an existing group. Rather, the most useful switches are those that let you modify the membership of a group, specifically

- **-addmbr *Member...*** adds members to the group specified in Group
- **-rmmbr *Member...*** removes members from the group specified in Group

where, as with all directory service commands, the DN is the full, distinguished name of another Active Directory object, surrounded by quotes if there are any spaces in the DN.

Note On any one command line, you can use only -addmbr *or* -rmmbr. You cannot use both in a single DSMOD GROUP command.

Practice: Using LDIFDE to Manage Group Accounts

In the following exercises, you list the options available for LDIFDE, export users from the Active Directory, and create a group object in the directory.

Exercise 1: Starting LDIFDE

In this exercise, you list the command options available with LDIFDE.

1. Open a Command Prompt.
2. For a list of commands, at the command prompt, type: **ldifde /?**.

Exercise 2: Exporting the Users from an Organizational Unit

In this exercise, you will export the entire contents of an OU named Marketing, complete with all its users, from the *contoso.com* domain.

1. In the *contoso.com* domain (Server01 is a domain controller for *contoso.com*), create an OU named Marketing.
2. In the Marketing OU, add two or three users. These users may be named whatever you choose.
3. Open a command prompt and type the following LDIFDE command (the character **:** indicates continuation to the next line)

```
ldifde -f marketing.ldf -s server01 :
-d "ou=Marketing,dc=contoso,dc=com" :
-p subtree -r : "(objectCategory=CN=Person,CN=Schema,CN=Configuration,:
DC=contoso,DC=com)"
```

Figure 4-4 shows the code in action.

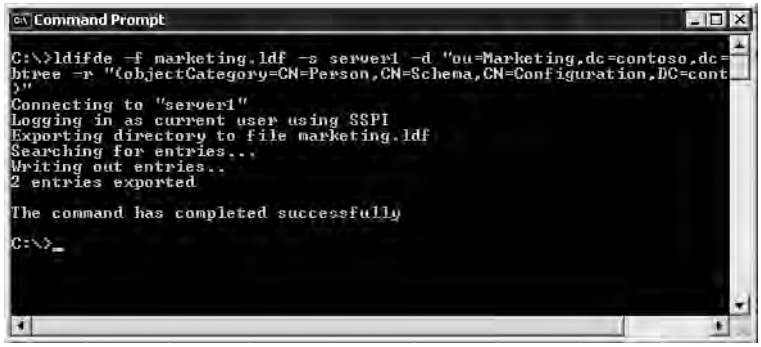

Figure 4-4 Output of LDIFDE export–Marketing OU

This creates a LDIF file named Marketing.ldf by connecting to the server named Server01 and executing a subtree search of the Marketing OU for all objects of the category Person.

Exercise 3: Using LDIFDE to Create a Group

In this exercise, you will use LDIFDE to add a group named Management to the Marketing OU of *contoso.com*.

1. Start a text editor, such as Notepad, and create a text file named Newgroup.ldf. (Save the file as an LDIF file, not as a text file.)
2. Edit the LDIF file Newgroup.ldf, and add the following text:

```
dn: CN=Management,OU=Marketing,DC=contoso,DC=com
changetype: add
cn: Management
objectClass: group
samAccountName: Marketing
```

3. Save and close the LDIF file.
4. Open a Command Prompt, type the following command and then press Enter:

   ldifde -i -f newgroup.ldf -s server01

Tip Watch for extra "white space" (tabs, spaces, carriage returns, line feeds) in the file. Extra white space in the file will cause the command to fail.

5. To confirm that the new group has been created, check the Active Directory Users And Computers snap-in.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Which of the following LDIFDE commands changes the function of LDIFDE from export to import?
   a. -i
   b. -t
   c. -f
   d. -s
2. What object classes are possible to export and import using LDIFDE?

_____

_____

3. You have a database of users that is capable of exporting CSV files. Can you use such a file, or must you create an *.ldf file manually for importing?

Lesson Summary

- LDIFDE is an included tool with Windows Server 2003 that allows for the importing and exporting of data into and out of Active Directory.
- If you have an existing directory of user data, you can use LDIFDE to export the desired data for importing into the Active Directory, which is, generally, a more efficient process than creating each element individually by hand. CSV files are usable, so long as the data is correctly formatted, with all required elements included and in their proper order.
- LDIFDE can be copied from a Windows Server 2003 to a Windows 2000 or Windows XP desktop for use with an Active Directory.

Case Scenario Exercise

You are in the process of building your Active Directory, and have some user data from the Human Resources department that includes first and last name, address, and telephone number. Company policy states that the user logon name should be the combination of first name or initial and last name (for example, Ben Smith would be bsmith).

You have 500 users, 30 groups, and 10 OUs. In practical terms, what is the best way to get your Active Directory set up as quickly and easily as possible?

Although there is no absolutely correct answer, there are different levels of complexity to consider. A blending of methods is probably best, given the following considerations:

- The user data can be edited as needed, but those edits are minimal, and the users can be brought into Active Directory using LDIFDE.
- The OU construction can be part of the user construction, all from the same file, with minimal editing. For the OUs, use LDIFDE as well.
- The groups might be another matter. Because group membership is a multivalued attribute in Active Directory, group membership must be listed, uniquely, for each group as it is created. It would be very confusing to do that within a single file, and errors would be likely. A better approach is to do the group memberships individually.

Troubleshooting Lab

Creating individual objects (users, groups, and computers) in your Active Directory is a straightforward process, but finding objects and their associations after many objects have been created can present challenges. In a large, multiple-domain environment (or in a complicated smaller one), solving resource access problems can be difficult. For example, if Sarah can access some but not all of the resources that are intended for her, she might not have membership in the groups that have been assigned permissions to the resources.

If you have multiple domains with multiple OUs in each domain, and multiple, nested groups in each of those OUs, it could take a great deal of time to examine the membership of these many groups to determine whether the user has the appropriate membership. Active Directory Users And Computers would not be the best tool choice.

You will use the DSGET command to get a comprehensive listing of all groups of which a user is a member. For the purposes of this lab, the user Ben Smith in the *contoso.com* domain, the Users OU will be used.

1. Choose a user in your Active Directory to use as a test case for the steps that follow. If you do not have a construction that is to your liking, create a number of nested groups across several OUs, making the user a member of only some of the groups.
2. Open a command prompt.
3. Type the following command (substituting your selected user name and OU for Ben Smith):

   ```
   dsget user "CN=Ben Smith,CN=Users,DC=contoso,DC=com"
   -memberof -expand
   ```

   The complete listing of all groups of which the user is a member is displayed.

Chapter Summary

- Groups may be created within any OU within the Active Directory.
- There are two types of groups: security and distribution.
- There are three scopes of groups: domain local, global, and universal.
- Manual creation of groups is accomplished with the Active Directory Users And Computers MMC.
- Automated creation of groups is accomplished with the LDIFDE command-line tool.

- Directory Services Tools such as DSQUERY, DSGET, and DSMOD can be used to list, create, and modify groups and their membership.
- Group types can only be changed when the domain functional level is at least Windows 2000 native.
- Advanced group nesting is only possible when the domain functional level is at least Windows 2000 native.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- The types of groups and their available uses depending on the domain functional level
- The scope of groups and their various nesting constructions depending on the domain functional level
- The basic use of Active Directory Users And Computers in creating groups and modifying their membership
- The basic use of LDIFDE for exporting groups from one directory to another, and in creating groups
- The basic use of DSGET for listing complete group memberships for a user

Key Terms

Domain local group (scope) In mixed or interim domain functional level, these local groups are available only on domain controllers, not domainwide.

Global group (scope) A group that is available domainwide in any domain functional level.

Universal group (scope) A group that can be available domainwide in any functional level, but limited to distribution scope in Windows 2000 mixed and Windows Server 2003 interim domain functional levels.

Security group (type) Can have permissions assigned in an ACL.

Distribution group (type) Cannot have permissions assigned in an ACL.

5 Computer Accounts

Exam Objectives in this Chapter:

- Create and manage computer accounts in a Microsoft Active Directory directory service environment
- Troubleshoot computer accounts
  - Diagnose and resolve issues related to computer accounts by using the Active Directory Users and Computers snap-in of the Microsoft Management Console (MMC)
  - Reset a computer account

Why This Chapter Matters

As an administrator, you are aware that, over time, hardware is added to your organization, computers are taken offline for repair, machines are exchanged between users or roles, and old equipment is retired or upgraded, leading to the acquisition of replacement systems. Each of these activities involves updating the computer accounts in Active Directory.

Just as a user is authenticated by the user object's user name and password, a computer maintains an account with a name and password that is used to create a secure relationship between the computer and the domain. A user can forget his or her password, requiring you to reset the password, or can take a leave of absence, requiring the disabling of the user object. Likewise, a computer's account can require reset or disabling.

In this chapter, you will learn how to create computer objects, which include the security properties required for the object to be an "account," and manage those objects using Active Directory Users And Computers' graphical user interface (GUI) as well as the powerful command-line tools of Microsoft Windows Server 2003. You will also review your understanding of the process through which a computer joins a domain, so that you can identify potential points of failure and more effectively troubleshoot computer accounts. Finally, you will master the key skills required to troubleshoot and repair computer accounts.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to computer accounts in Active Directory. If you desire hands-on practice, using the examples and lab exercises in the chapter, you should have the following prepared:

- A machine running Windows Server 2003 (Standard Edition or Enterprise Edition) installed as Server01 and configured as a domain controller in the domain *contoso.com*.
- First-level organizational units (OUs): "Administrative Groups," "Desktops," and "Servers."
- A global security group, in the Administrative Groups OU, called "Deployment."
- The Active Directory Users And Computers console, or a customized console with the Active Directory Users And Computers snap-in.
- One exercise, joining a computer to a domain, is possible only if you have a second computer running Microsoft Windows 2000 Professional, Windows XP, or Windows Server 2003, with connectivity to Server01. DNS services must be configured properly, on Server01 or elsewhere, and the second computer must be configured to use that DNS server, so that it can locate the domain controller (Server01) for *contoso.com*.

Lesson 1: Joining a Computer to a Domain

The default configuration of Windows Server 2003, and all Microsoft Windows operating systems, is that the computer belongs to a workgroup. In a workgroup, a Windows NT–based computer (which includes Windows NT 4, Windows 2000, Windows XP, and Windows Server 2003) can authenticate users only from its local Security Accounts Manager (SAM) database. It is a stand-alone system, for all intents and purposes. Its workgroup membership plays only a minor role, specifically in the browser service. Although a user at that computer can connect to shares on other machines in a workgroup or in a domain, the user is never actually logged on to the computer with a domain account.

Before you can log on to a computer with your domain user account, that computer must belong to a domain. The two steps necessary to join a computer to a domain are, first, to create an account for the computer and, second, to configure the computer to join the domain using that account. This lesson will focus on the skills related to the creation of computer accounts and joining computers to domains. The next lesson will explore, in more depth, the computer accounts themselves.

Computers maintain accounts, just as users do, that include a name, password, and security identifier (SID). Those properties are incorporated into the computer object class within Active Directory. Preparing for a computer to be part of your domain is therefore a process strikingly similar to preparing for a user to be part of your domain: you must create a computer object in Active Directory.

After this lesson, you will be able to

- Create computer accounts using Active Directory Users And Computers
- Create computer accounts using the DSADD command-line tool
- Create computer accounts using the NETDOM command-line tool
- Join a computer to a domain by changing the network identification properties
- Understand the importance of creating computer accounts prior to joining a domain

Estimated lesson time: 20 minutes

Creating Computer Accounts

You must be a member of the Administrators or Account Operators groups on the domain controllers to create a computer object in Active Directory. Domain Admins and Enterprise Admins are, by default, members of the Administrators group. Alternatively, it is possible to delegate administration so that other users or groups can create computer objects.

However, domain users can also create computer objects through an interesting, indirect process. When a computer is joined to the domain and an account does not exist, Active Directory creates a computer object automatically, by default, in the Computers OU. Each user in the Authenticated Users group (which is, in effect, all users) is allowed to join 10 computers to the domain, and can therefore create as many as 10 computer objects in this manner.

Creating Computer Objects Using Active Directory Users and Computers

To create a computer object, or "account," open Active Directory Users And Computers and select the container or OU in which you want to create the object. From the Action menu or the right-click shortcut menu, choose the New–Computer command. The New Object–Computer dialog box appears, as illustrated in Figure 5-1.

Figure 5-1 The New Object–Computer dialog box

In the New Object–Computer dialog box, type the computer name. Other properties in this dialog box will be discussed in the following lesson. Click Next. The following page of the dialog box requests a GUID. A GUID is used to prestage a computer account for Remote Installation Services (RIS) deployment, which is beyond the scope of this discussion. It is not necessary to enter a GUID when creating a computer account for a machine you will be joining to the domain using other methods. So just click Next and then click Finish.

Creating Computer Objects Using DSADD

Chances are, this is something you've done before. But before you decide there's nothing new under the sun, Windows Server 2003 provides a useful command-line tool, DSADD, which allows you to create computer objects from the command prompt or a batch file.

In Chapter 2, "Administering Microsoft Windows Server 2003," you used DSADD to create user objects. To create computer objects, simply type **dsadd computer ComputerDN**, where ComputerDN is the distinguished name (DN) of the computer, such as CN=Desktop123,OU=Desktops,DC=contoso,DC=com.

If the computer's DN includes a space, surround the entire DN with quotation marks. The *ComputerDN…* parameter can include more than one distinguished name for new computer objects, making DSADD Computer a handy way to generate multiple objects at once. The parameter can be entered in one of the following ways:

- By piping a list of DNs from another command, such as dsquery.
- By typing each DN on the command line, separated by spaces.
- By leaving the DN parameter empty, at which point you can type the DNs, one at a time, at the keyboard console of the command prompt. Press ENTER after each DN. Press CTRL+Z and ENTER after the last DN.

The DSADD Computer command can take the following optional parameters after the DN parameter:

- -samid *SAMName*
- -desc *Description*
- -loc *Location*

Creating a Computer Account with NETDOM

The NETDOM command is available as a component of the Support Tools, installable from the Support\Tools directory of the Windows Server 2003 CD. The command is also available on the Windows XP and Windows 2000 CDs. Use the version that is appropriate for the platform. NETDOM allows you to perform numerous domain account and security tasks from the command line.

To create a computer account in a domain, type the following command:

```
netdom add ComputerName /domain:DomainName /userd:User /PasswordD:UserPassword
[/ou:OUDN]
```

This command creates the computer account for *ComputerName* in the domain *DomainName* using the domain credentials *User* and *UserPassword*. The /ou parameter causes the object to be created in the OU specified by the *OUDN* distinguished name following the parameter. If no OUDN is supplied, the computer account is created in the Computers OU by default. The user credentials must, of course, have permissions to create computer objects.

Joining a Computer to a Domain

A computer account alone is not enough to create the secure relationship required between a domain and a machine. The machine must join the domain.

To join a computer to the domain, perform the following steps:

1. Right-click My Computer and choose Properties. Click the Computer Name tab.
   - ❑ Open Control Panel, select System, and in the System Properties dialog box, click the Computer Name tab.
   - ❑ Open the computer's Computer Name properties. These properties can be accessed in several ways:

Note The Computer Name tab is called Network Identification on Windows 2000 systems. The Change button is called Properties. The functionality is, however, identical.

2. Open the Network Connections folder from Control Panel and choose the Network Identification command from the Advanced menu.
3. On the Computer Name tab, click Change. The Computer Name Changes dialog box, shown in Figure 5-2 allows you to change the name and the domain and workgroup membership of the computer.

Exam Tip You will not be able to change a computer's name or membership if you are not logged on with administrative credentials on that system. Only users who belong to the local Administrators group will find the Change button enabled and functional.

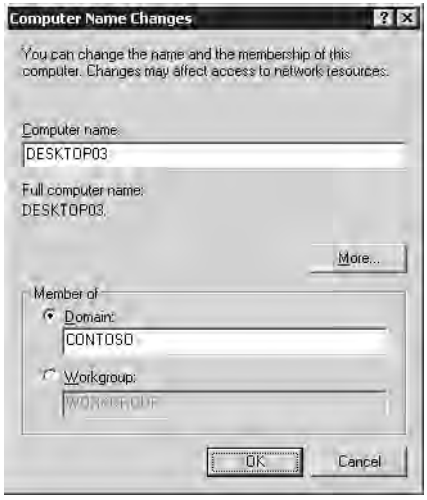

Figure 5-2 The Computer Name Changes dialog box

4. In the Computer Name Changes dialog box, click Domain and type the name of the domain.

Tip Although the NetBIOS (flat) domain name may succeed in locating the target domain, it is best practice to enter the DNS name of the target domain. DNS configuration is critical to a Windows 2000, Windows XP, or Windows Server 2003 computer. By using the DNS domain name, you leverage the preferred name resolution process and test the computer's DNS configuration. If the computer is unable to locate the domain you're attempting to join, ensure that the DNS server entries configured for the network connection are correct.

5. Click OK. The computer contacts the domain controller. If there is a problem connecting to the domain, examine network connectivity and configuration, as well as DNS configuration.

When the computer successfully contacts the domain, you will be prompted, as in Figure 5-3, for a user name and password with privileges to join the domain. Note that the credentials requested are your *domain* user name and password.

Figure 5-3 Prompt for credentials to join domain

If you have *not* created a domain computer account with a name that matches the computer's name, Active Directory creates an account automatically in the default Computers container. Once a domain computer account has been created or located, the computer establishes a trust relationship with the domain, alters its SID to match that of the account, and makes modifications to its group memberships. The computer must then be restarted to complete the process.

Note The NETDOM JOIN command can also be used to join a workstation or server to a domain. Its functionality is identical to the Computer Name Changes user interface, except that it also allows you to specify the OU in which to create an account if a computer object does not already exist in Active Directory.

The Computers Container vs. OUs

The Computers container is the default location for computer objects in Active Directory. After a domain is upgraded from Windows NT 4 to Windows 2000, all computer accounts are found, initially, in this container. Moreover, when a machine joins the domain and there is no existing account in the domain for that computer, a computer object is created automatically in the Computers container.

Tip The *Microsoft Windows Server 2003 Resource Kit* includes the REDIRCOMP tool, which allows you to redirect the creation of automatic computer objects to an OU of your choice. The domain must be in Windows Server 2003 Domain functionality, meaning that all domain controllers must be running Windows Server 2003. Such a tool is useful to organizations in which computer account creation is less tightly controlled. Because automatically created computer objects are created in an OU, they can be managed by policies linked to that OU. See the *Windows Server 2003 Resource Kit* for more information on REDIRCOMP.

Although the Computers container is the default container for computer objects, it is not the ideal container for computer objects. Unlike OUs, containers such as Computers, Users and Builtin cannot be linked to policies, limiting the possible scope of computer-focused group policy. A best-practice Active Directory design will include at least one OU for computers. Often, there are multiple OUs for computers, based on administrative division, region, or for the separate administration of laptops, desktops, file and print servers, and application servers. As an example, there is a default OU for Domain Controllers in Active Directory, which is linked to the Default Domain Controller Policy. By creating one or more OUs for computers, an organization can delegate administration and manage computer configuration, through group policy, more flexibly.

If your organization has one or more OUs for computers, you must move any computer objects created automatically in the Computers container into the appropriate OU. To move a computer object, select the computer and choose Move from the Action menu. Alternatively, use the new drag-and-drop feature of the MMC to move the object.

Tip Because a computer object in the Computers OU will not be governed by the group policies linked to the OUs your organization has created specifically for computers; and because it requires an extra step to move a computer object from the Computers OU into the appropriate OU, it is recommended to *create computer objects before joining the computer to the domain*. You can create the computer object in the correct OU initially, so that once the system joins the domain it is immediately governed by the policies linked to that OU.

You can also move a computer object, or any other object, with the DSMOVE command. The syntax of DSMOVE is:

```
dsmove ObjectDN [-newname NewName] [-newparent ParentDN]
```

The -newname parameter allows you to rename an object. The -newparent parameter allows you to move an object. To move a computer named DesktopABC from the Computers container to the Desktops OU, you would type the following:

```
dsmove ?CN=DesktopABC,CN=Computers,DC=Contoso,DC=com? -newparent
?OU=Desktops,DC=Contoso,DC=com?
```

In this command you again see the distinction between the Computers *container* (CN) and the Desktops *organizational unit* (OU).

You must have appropriate permissions to move an object in Active Directory. Default permissions allow Account Operators to move computer objects between containers including the Computers container and any OUs *except* into or out of the Domain Controllers OU. Administrators, which include Domain Admins and Enterprise Admins, can move computer objects between any containers, including the Computers container, the Domain Controllers OU, and any other OUs.

Practice: Joining a Computer to an Active Directory Domain

In this practice, you will create computer accounts using Active Directory Users and Computers and DSADD. You then can join a computer to the domain, if you have access to a second system.

Exercise 1: Creating Computer Accounts with Active Directory Users and Computers

1. Open Active Directory Users And Computers
2. In the Servers OU, create a computer object for a computer named "SERVER02." Configure only the computer name. Do not change any of the other default properties.

   Note that, like a user, a computer has two names—the computer name and the "Pre–Windows 2000" computer name. It is a best practice to keep the names the same.

Exercise 2: Creating Computer Accounts with DSADD

1. Open the command prompt.
2. Type the command:

   ```
   dsadd computer ?cn=desktop03,ou=servers,dc=contoso,dc=com?
   ```

Exercise 3: Moving a Computer Object

1. Open Active Directory Users And Computers.
2. Using the Move command, move the Desktop03 computer object from the Servers OU to the Desktops OU.
3. Drag Server02 from the Servers container to the Computers container.
4. Select the Computers container to confirm that Server02 arrived in the right place. Drag-and-drop is, of course, subject to user error.

Off the Record The MMC is notorious for causing mild panic attacks. It does *not* refresh automatically. You must use the Refresh command or shortcut key (F5) to refresh the console after making a change such as moving an object.

5. Open the properties of the Computers container. You will see that it does *not* have a Group Policy tab, unlike an OU such as Servers. This is among the reasons why organizations create one or more additional OUs for computer objects.
6. Open a command prompt.
7. Type the command:

```
dsmove ?CN=Server02,CN=Computers,DC=contoso,DC=com? -newparent
?OU=Servers,DC=contoso,DC=com?
```

   This command, as you can deduce, will move the computer object back to the Servers OU.

8. Confirm that the computer is again in the Servers OU.

Exercise 4 (Optional): Join a Computer to a Domain

This exercise requires an additional system with network connectivity to Server01. In addition, DNS must be configured correctly so that Server01's service records (SRV) are created. The additional computer must have DNS configured so that it can locate Server01 as a domain controller for *contoso.com*.

1. If you have an additional system that you are able to join to the domain in the next exercise, create an account for it in the Desktops OU using either Active Directory Users And Computers or DSADD. Be certain that the name you use is the same name as the computer.
2. Log on to the computer. You must log on as an account with membership in the computer's local Administrators group to change its domain membership.

3. Locate the Computer Name tab by opening System from Control Panel, or the Network Identification command from the Advanced menu of the Network Connections folder.
4. Click Change.
5. Click Domain and type the DNS domain name, **contoso.com**.
6. Click OK.
7. When prompted, enter the credentials for the *contoso.com* domain's Administrator account.
8. Click OK.
9. The computer will prompt you that a reboot is necessary. Click OK to each message and to close each dialog box. Reboot the system.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What are the *minimum* credentials necessary to create a Windows Server 2003 computer account in an OU in a domain? Consider all steps of the process. Assume Active Directory does not yet have an account for the computer.
   a. Domain Admins
   b. Enterprise Admins
   c. Administrators on a domain controller
   d. Account Operators on a domain controller
   e. Server Operators on a domain controller
   f. Account Operators on the server
   g. Server Operators on the server
   h. Administrators on the server
2. Which locations allow you to change the domain membership of a Windows Server 2003 computer?
   a. The properties of My Computer
   b. Control Panel's System application
   c. Active Directory Users and Computers
   d. The Network Connections folder
   e. The Users application in Control Panel

3. What command-line tools will create a domain computer account in Active Directory?

   a. NETDOM

   b. DSADD

   c. DSGET

   d. NETSH

   e. NSLOOKUP

Lesson Summary

- Members of the Administrators and Account Operators groups have, by default, permission to create computer objects in Active Directory.
- Active Directory Users And Computers, DSADD, and NETDOM can be used to create computer accounts.
- You must be logged on as a member of the *local* Administrators group to change the domain membership of a machine.

Lesson 2: Managing Computer Accounts

In the previous lesson, you examined the fundamental components of a computer's relationship with a domain: the computer's account, and joining the computer to the domain. This lesson looks more closely at the computer object in Active Directory. You will learn about the other properties and permissions that make computer objects "tick," and how to manage those properties and permissions using GUI and command-line tools.

After this lesson, you will be able to

- Configure the permissions of a new Active Directory computer object
- Configure the properties of an Active Directory computer object
- Find and manage computer accounts using Active Directory Users And Computers

Estimated lesson time: 10 minutes

Managing Computer Object Permissions

In Lesson 1, you learned that you could join a computer to a domain by providing domain administrator credentials when prompted by the computer during the join process. Security concerns, however, require us to use the minimum necessary credentials to achieve a particular task, and it does seem like overkill to need a Domain Admins' account to add a desktop to the domain.

Fortunately, Active Directory allows you to control, with great specificity, the groups or users that can join a computer to a domain computer account. Although the default is Domain Admins, you can allow any group (for example, a group called "Installers") to join a machine to an account. This is most easily achieved while creating the computer object.

When you create a computer object, the first page of the New Object–Computer dialog box (previously shown in Figure 5-1) indicates The Following User Or Group Can Join This Computer To A Domain. Click Change and you can select any user or group. This change modifies a number of permissions on the computer object in Active Directory.

The following page of the New Object–Computer dialog box prompts you for the globally unique identifier (GUID) of the computer, which is necessary if you install a system using Remote Installation Services (RIS). For more information on RIS, see the Microsoft online Knowledge Base, *http://support.microsoft.com/*.

If the computer that is using the account that you are creating is running a version of Windows earlier than 2000, select the Assign This Computer Account As A Pre–Windows 2000 Computer check box. If the account is for a Windows NT backup domain controller, click Assign This Computer Account As A Backup Domain Controller.

Tip Remember, only computers based on Windows NT technologies can belong to a domain, so Windows 95, Windows 98, and Windows Millennium Edition (Windows Me) cannot join or maintain computer accounts. Therefore, this check box really means Windows NT 4.

Configuring Computer Properties

Computer objects have several properties that are not visible when creating a computer account in the user interface. Open a computer object's Properties dialog box to set its location and description, configure its group memberships and dial-in permissions, and link it to a user object of the computer's manager. The Operating System properties page is read-only. The information is published automatically to Active Directory, and will be blank until a computer has joined the domain using that account.

Several object classes in Active Directory support the Manager property that is shown on the Managed By property page of a computer. This linked property creates a cross-reference to a user object. All other properties—the addresses and telephone numbers—are displayed directly from the user object. They are not stored as part of the computer object itself.

The DSMOD command, as discussed in Chapter 2, can also modify several of the properties of a computer object. You will see the DSMOD command in action in the following section regarding troubleshooting computer accounts.

Finding and Connecting to Objects in Active Directory

When a user calls you with a particular problem, you might want to know what operating system and service pack is installed on that user's system. You learned that this information is stored as properties of the computer object. The only challenge, then, is to locate the computer object, which may be more difficult in a complex Active Directory with one or more domains and multiple OUs.

The Active Directory Users and Computers snap-in provides easy access to a powerful, graphical search tool. This tool can be used to find a variety of object types. In this context, however, your search entails an object of the type Computer. Click the Find Objects In Active Directory button on the console toolbar. The resulting Find Computers dialog box is illustrated in Figure 5-4. You can select the type of object (Find), the scope of the search (In), and specify search criteria before clicking Find Now.

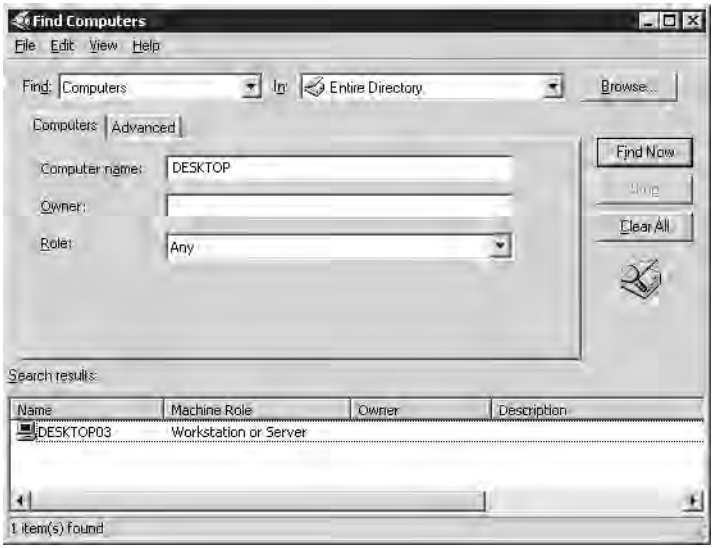

Figure 5-4 The Find Computers dialog box, as it appears after a successful search

The list of results allows you to select an object and, from the File menu or the shortcut menu, perform common tasks on the selected object. Many administrators appreciate learning that you can use the Manage command to open the Computer Management console and connect directly to that computer, allowing you to examine its event logs, device manager, system information, disk and service configuration, or local user or group accounts.

Practice: Managing Computer Accounts

In this practice, you will search for a computer object and modify its properties.

Exercise 1: Managing Computer Accounts

1. Open Active Directory Users And Computers.
2. Select the Security Groups OU and create a global security group called Deployment.
3. Select the Desktops OU.
4. Create a computer account for Desktop04. In the first page of the New Object–Computer dialog box, click Change below The Following User Or Group Can Join This Computer To A Domain. Type **deployment** in the Select User or Group dialog box, and then click OK.
5. Complete the creation of the Desktop04 computer object.

Exercise 2: Finding Objects in Active Directory

1. Open Active Directory Users And Computers.
2. On the toolbar, click the Find Objects in Active Directory icon.
3. By default, the Find dialog box is ready to search for Users, Contacts, and Groups. Choose Computers from the Find drop-down list, and select Entire Directory from the In drop-down list.
4. In the Computer Name field, type **server** and click Find Now.

   A result set appears that includes Server01.

Exercise 3: Changing Computer Properties

1. From the result set returned in Exercise 1, open Server01's properties dialog box.
2. Click the Location tab.
3. Type **Headquarters Server Room**.
4. Click the Managed By tab, and then click Change.
5. Type **Hank** and then click OK.
6. Note that the user's name and contact information appears.
7. Click the Operating System tab. Note the OS version and service pack level are displayed.
8. (Optional) If you joined a second computer to the domain in Exercise 4 of Lesson 1, open the properties of that computer object and note the Operating System properties of that computer.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What platforms are capable of joining a domain?
   a. Windows 95
   b. Windows NT 4
   c. Windows 98
   d. Windows 2000
   e. Windows Me

f. Windows XP

g. Windows Server 2003

2. You open a computer object and, on the Operating System tab, discover that no properties are displayed. What causes these properties to be absent?

3. An executive has a laptop running Windows XP, with a machine name of "TopDog." You want to allow the executive's laptop to join the domain, and you want to be sure that the computer is configured by the group policies linked to the Desktops OU immediately. How can you achieve this goal?

4. Why is it a best practice to create a computer account in the domain *prior* to joining a machine to the domain?

Lesson Summary

- You can allow any user or group to join a computer to a domain account by utilizing the property, The Following User Or Group Can Join This Computer To A Domain.
- The Find Objects In Active Directory button on the Active Directory Users And Computers snap-in toolbar allows you to search for, and then manage, computer and other Active Directory objects.

Lesson 3: Troubleshooting Computer Accounts

Active Directory domains treat computers as security principals. This means that a computer, just like a user, has an account—or, more specifically, properties within the computer object such as a name, a password, and a SID. Like user accounts, computer accounts require maintenance and, occasionally, troubleshooting. This lesson focuses on skills and concepts related to troubleshooting computer objects.

After this lesson, you will be able to

- Understand the important difference among deleting, disabling, and resetting computer accounts
- Recognize the symptoms of computer account problems
- Troubleshoot computer accounts by deleting, disabling, resetting, or rejoining, using both command-line and user-interface tools

Estimated lesson time: 20 minutes

Deleting and Disabling and Resetting Computer Accounts

Computer accounts, like user accounts, maintain a unique SID, which enables an administrator to grant permissions to computers. Also like user accounts, computers can belong to groups. Therefore, like user accounts, it is important to understand the effect of deleting a computer account. When a computer account is deleted, its group memberships and SID are lost. If the deletion is accidental, and another computer account is created with the same name, it is nonetheless a new account, with a new SID. Group memberships must be reestablished, and any permissions assigned to the deleted computer must be reassigned to the new account. Delete computer objects only when you are certain that you no longer require those security-related attributes of the object.

To delete a computer account using Active Directory Users And Computers, locate and select the computer object and, from the Action menu or the shortcut menu, select the Delete command. You will be prompted to confirm the deletion and, because deletion is not reversible, the default response to the prompt is No. Select Yes and the object is deleted.

The DSRM command-line tool introduced in Chapter 3 allows you to delete a computer object from the command prompt. To delete a computer with DSRM, type:

```
DSRM ObjectDN
```

Where *ObjectDN* is the distinguished name of the computer, such as "CN=Desktop15, OU=Desktops,DC=contoso,DC=com." Again, you will be prompted to confirm the deletion.

Tip When a computer is disjoined from a domain—when an administrator changes the membership of the computer to a workgroup or to another domain—the computer attempts to delete its computer account in the domain. If it is not possible to do so because of lack of connectivity, networking problems, or credentials and permissions, the account will remain in Active Directory. It may appear, immediately or eventually, as disabled. If that account is no longer necessary, it must be deleted manually.

If a computer is taken offline or is not to be used for an extended period of time, you may disable the account. Such an action reflects the security principle, that an identity store allow authentication only of the minimum number of accounts required to achieve the goals of an organization. Disabling the account does not modify the computer's SID or group membership, so when the computer is brought back online, the account can be enabled.

The context menu, or Action menu, of a selected computer object exposes the Disable Account command. A disabled account appears with a red "X" icon in the Active Directory Users And Computers snap-in, as shown in Figure 5-5.

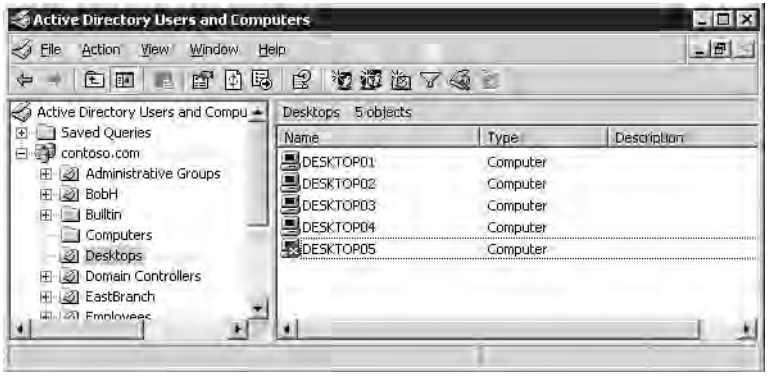

Figure 5-5 A disabled computer account

While an account is disabled, the computer cannot create a secure channel with the domain. The result is that users who have not previously logged on to the computer, and who therefore do not have cached credentials on the computer, will be unable to log on until the secure channel is reestablished by enabling the account.

To enable a computer account, simply select the computer and choose the Enable Account command from the Action or shortcut menus.

To disable or enable a computer from the command prompt, use the DSMOD command. The DSMOD command modifies Active Directory objects. The syntax used to disable or enable computers is:

```
DSMOD COMPUTER ComputerDN -DISABLED YES

DSMOD COMPUTER ComputerDN -DISABLED NO
```

If a computer account's group memberships and SID, and the permissions assigned to that SID, are important to the operations of a domain, you do not want to delete that account. So what would you do if a computer was replaced with a new system, with upgraded hardware? Such is one scenario in which you would *reset* a computer account.

Resetting a computer account resets its password, but maintains all of the computer object's properties. With a reset password, the account becomes in effect "available" for use. Any computer can then join the domain using that account, including the upgraded system.

In fact, the computer that had previously joined the domain with that account can use the reset account by simply rejoining the domain. This reality will be explored in more detail in the troubleshooting lesson.

The Reset Account command is available in the Action and context menus when a computer object is selected. The DSMOD command can also be used to reset a computer account, with the following syntax:

```
dsmod computer ComputerDN -reset
```

The NETDOM command, included with the Windows Server 2003 Support Tools in the CD-ROM's Support\Tools directory, also enables you to reset a computer account.

Recognizing Computer Account Problems

Computer accounts, and the secure relationships between computers and their domain are robust. In the rare circumstance that an account or secure channel breaks down, the symptoms of failure are generally obvious. The most common signs of computer account problems are:

Messages at logon indicate that a domain controller cannot be contacted; that the computer account may be missing; or that the trust (another way of saying "the secure relationship") between the computer and the domain has been lost. An example is shown in Figure 5-6.

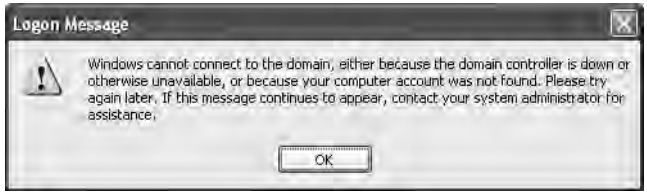

Figure 5-6 Logon message from a Windows XP client indicating a possible computer account problem

- Error messages or events in the event log indicating similar problems or suggesting that passwords, trusts, secure channels, or relationships with the domain or a domain controller have failed.
- A computer account is missing in Active Directory.

If one of these situations occurs, you must troubleshoot the account. You learned earlier how to delete, disable, and reset a computer account and, at the beginning of the chapter, how to join a machine to the domain.

The rules that govern troubleshooting a computer account are:

A. If the computer account exists in Active Directory, it must be reset.

B. If the computer account is missing in Active Directory, you must create a computer account.

C. If the computer still belongs to the domain, it must be removed from the domain by changing its membership to a workgroup. The name of the workgroup is irrelevant. Best practice is to try and choose a workgroup name that you know is not in use.

D. Rejoin the computer to the domain. Alternatively, join another computer to the domain; but the new computer must have the same name as the computer account.

To troubleshoot any computer account problem, apply *all four rules*. These rules can be addressed in any order, except that Rule D, involving rejoining the computer to the domain, must as always be performed as the final step. Let's examine two scenarios.

In the first scenario, a user complains that when he or she attempts to log on, the system presents error messages indicating that the computer account might be missing. Applying Rule A, you open Active Directory Users And Computers and find that the computer account exists. You reset the account. Rule B does not apply—the account does exist. Then, using Rule C, you disjoin the system from the domain and, following Rule D, rejoin the domain.

In a second scenario, if a computer account is reset by accident, the first item that has occurred is Rule A. Although the reset is accidental, you must continue to recover by

applying the remaining three rules. Rule B does not apply because the account exists in the domain. Rule C indicates that if the computer is still joined to the domain, it must be removed from the domain. Then, by Rule D, it can rejoin the domain.

With these four rules, you can make an informed decision, on the job or on the certification exams, about how to address any scenario in which a computer account has lost functionality.

Practice: Troubleshooting Computer Accounts

In this practice, you will troubleshoot a realistic scenario. A user in the *contoso.com* domain contacts you and complains that, when logging on to Desktop03, he or she receives the following error message:

"Windows cannot connect to the domain, either because the domain controller is down or otherwise unavailable, or because your computer account was not found. Please try again later. If this message continues to appear, contact your system administrator for assistance."

The user waited, attempted to log on, received the same message, waited again, and then received the same message a third time. The user has now spent 20 minutes trying to log on. In obvious frustration, the user contacts you for assistance.

Exercise 1: Troubleshooting Computer Accounts

1. Identify the most likely cause of the user's problem:
   a. The user entered an invalid user name.
   b. The user entered an invalid password.
   c. The user chose the incorrect domain from the Log On To list.
   d. The computer has lost its secure channel with the domain.
   e. The computer's registry is corrupted.
   f. The computer has a policy preventing the user from logging on interactively.

   The correct answer, as you can probably deduce, is d. The computer has lost its secure channel with the domain.

2. Identify the steps from the list below that you must take to troubleshoot the problem. Put the steps in order. You may not require all steps.
   a. Enable the computer account.
   b. Change Desktop03 to belong to *contoso.com*.
   c. Determine whether the computer account exists in Active Directory.
   d. Reset or re-create the computer account.

e. Change Desktop03 to a workgroup.

f. Delete the computer account.

g. Disable the computer account.

The correct answer is steps e, c, d, and b. Step e does not have to occur first; it just has to be done anytime before step b. Steps c and d must occur, in that order, before step b, which must be the last step.

Exercise 2: Recover from Computer Account Problems

1. Open Active Directory Users And Computers.
2. Click Find Objects In Active Directory and search for Desktop03.
3. Desktop03 appears in the search results because you created it in Lesson 1.
4. Having identified that the account does exist, reset the account by right-clicking Desktop03 and choosing Reset Account.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. After a period of expansion, your company created a second domain. Last weekend, a number of machines that had been in your domain were moved to the new domain. When you open Active Directory Users And Computers, the objects for those machines are still in your domain, and are displayed with a red "X" icon. What is the most appropriate course of action?

   a. Enable the accounts

   b. Disable the accounts

   c. Reset the accounts

   d. Delete the accounts

2. A user reports that during a logon attempt, a message indicated that the computer cannot contact the domain because the domain controller is down or the computer account may be missing. You open Active Directory Users And Computers and discover that the account for that computer is missing. What steps should you take?

   _____

   _____

3. A user reports that during a logon attempt, a message indicates that the computer cannot contact the domain because the domain controller is down or the computer account may be missing. You open Active Directory Users and Computers and that computer's account appears normal. What steps should you take?

Lesson Summary

- Computers maintain accounts that, like users, include a SID and group memberships. Be careful about deleting computer objects. Disabling computer objects allows you to enable the objects again, when the computer needs to participate in the domain.
- Problems with computer accounts are generally quite evident, with error messages and events logged that indicate problems in an account, a password, a secure channel or a trust relationship.
- Using the four rules in Lesson 3, you can troubleshoot just about any computer account problem.

Case Scenario Exercise

Contoso decides to open two branch offices: East and West. Computers are purchased for 10 sales representatives in each office. The asset tags assigned to the computers are shown in the following table.

East Branch	West Branch
EB-2841	WB-3748
EB-2842	WB-3749
EB-2843	WB-3750
EB-2844	WB-3751
EB-2845	WB-3752
EB-2846	WB-3753
EB-2847	WB-3754
EB-2848	WB-3755
EB-2849	WB-3756
EB-2850	WB-3757

Your job is to prepare Active Directory for the deployment of these computers.

Exercise 1: Create OUs

Create two OUs in the *contoso.com* domain: EastBranch and WestBranch. Type the names as shown. Do not put a space between the words.

Exercise 2: Script the Creation of Computer Accounts

1. Open Notepad.
2. Type a line for each computer, following this example:

   ```
   DSADD COMPUTER ?CN=EB-2841,OU=EastBranch,DC=Contoso,DC=COM? -desc ?Sales Rep
   Computer? -loc ?East Branch Office?
   ```

   Be sure to modify the CN= parameter to match the asset tag of each computer, and the OU= and -loc parameters to reflect the name and location description of the branch office for each computer.
3. Save the file as "C:\ScriptComputers.bat" and be sure to surround the name with quotation marks, or Notepad will add a .txt extension automatically.
4. Open a command prompt and type **c:\scriptcomputers**.
5. Confirm the successful generation of the computer accounts by examining the EastBranch and WestBranch OUs. The MMC does not refresh automatically, so press F5 to refresh if you do not see the new computers initially.

Troubleshooting Lab

Following a weekend during which a consultant performed maintenance on the computers in the East Branch Office, users complain of trouble logging on. You examine the event log on one of the branch office computers and discover the following event:

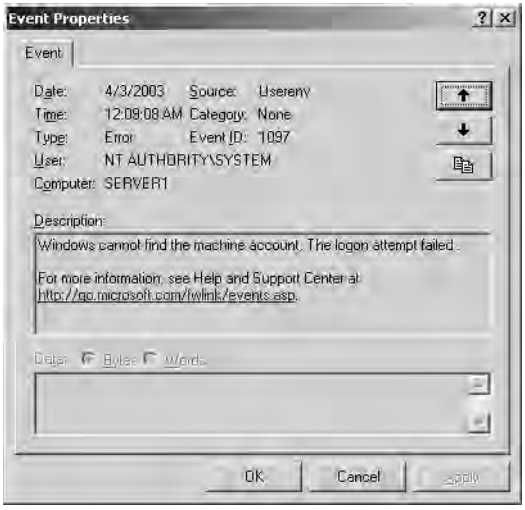

There seems to be a problem with the computer account.

Which of the following steps must be performed to correct the problem?

1. Delete the computer accounts
2. Reset the user accounts
3. Join the computers to a workgroup
4. Disable the computer accounts
5. Reset the computer accounts
6. Enable the computer accounts
7. Create new computer accounts
8. Join the computers to the domain

The correct answer is 5, 3, and 8. This is the most efficient solution; it involves resetting computer accounts and rejoining machines to the domain.

Exercise 1 (Optional): Simulation of the Problem

If you joined a second computer to the Contoso domain in Lesson 1, move the computer object for that computer into the EastBranch OU. Then, in Active Directory Users And Computers, reset the computer's account.

When you restart the computer, try logging on to the domain. Are you successful? Can you log on with Contoso domain accounts you have used in the past to log on to the computer? Why? (Hint: cached logons.)

Can you log on with new domain accounts, which have never logged on to the computer? When you attempt to do so, you will receive a typical error message indicating that the computer account may be missing.

Log on as the local Administrator and examine the event log. What error messages appear?

Exercise 2: Reset All East Branch Computer Accounts

The fastest way to reset the computer accounts, particularly because all the accounts are in the same OU, will be a command-line tool.

1. Open a command prompt.
2. Type the following command:

```
DSQUERY COMPUTER ?OU=EastBranch,DC=contoso,DC=com?
```

This command queries Active Directory for a list of computers in the EastBranch OU. The list should match the computer accounts created in the Case Scenario exercise.

3. Type the following command:

```
DSQUERY COMPUTER ?OU=EastBranch,DC=contoso,DC=com? | DSMOD COMPUTER -RESET
```

This time, we pipe the results of the DSQUERY command to the input of DSMOD. The DSMOD COMPUTER -RESET command will reset each of those accounts. Mission accomplished.

Exercise 3 (Optional): Rejoin the Domain

If you have a second system, just reset its computer account. You can now practice removing the machine from the domain by changing its membership to a workgroup. After restarting, join the domain again.

Chapter Summary

- You must have permissions to create a computer object in Active Directory. Administrators and Account Operators have sufficient permissions, and permissions can be delegated to other users or groups.
- When creating a computer object, you can specify what user or group can join the computer to the domain using that account.
- Active Directory Users And Computers allows you to create, modify, delete, disable, enable, and reset computer objects.
- From the command prompt, you can create a computer object with DSADD Computer and modify its properties using DSMOD Computer.
- DSMOD Computer is also used to reset, disable, and enable a computer object. DSRM will remove a computer object. The support tool, NETDOM, includes numerous switches to achieve similar tasks.
- A common troubleshooting recovery includes re-creating or resetting a computer account, removing the computer from the domain, and rejoining the domain.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the “Further Readings” sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Identify the minimum permissions required to create a computer object in Active Directory, and the permissions required to change a machine's membership between workgroups and domains.
- Know the syntax of the DSADD, DSMOD, and DSRM commands. Remember that DSMOD and DSADD require one, or more, distinguished names as parameters. The DSQUERY command can be used to provide those names to DSMOD.
- Be very clear on the differences among disabling, resetting, and deleting a computer account. What is the impact of each on the computer object, its SID and group membership, and on the system itself?
- Know the four rules for troubleshooting computer account problems. Apply all four, every time, and you will be likely to nail every computer account troubleshooting question.
- Be comfortable with finding objects in Active Directory, and managing those objects from the search results. This skill set applies to many objects in Active Directory, and several objectives of the certification exam.

Key Terms

Computer account An account created in Active Directory that uniquely identifies the computer in the domain.

6 Files and Folders

Exam Objectives in this Chapter:

- Configure access to shared folders
  - Manage shared folder permissions
- Configure file system permissions
  - Verify effective permissions when granting permissions
  - Change ownership of files or folders
- Troubleshoot issues related to access to files and shared folders
- Manage a Web server
  - Manage Internet Information Services (IIS)
  - Manage security for IIS

Why This Chapter Matters

Among the more common daily challenges facing you as an administrator are tasks related to the maintenance of network files and folders—resources that are required by users in your organization. When a user cannot access a resource that he or she needs to achieve a business task, the telephone at the help desk rings. As a result, you spend time and money modifying permissions or group memberships to correct the problem. When a sensitive resource is accessed by someone who should not be able to do so, the telephone on your desk rings—and as a result, you might have to spend time and money looking for a new job.

You have no doubt experienced the fundamental components of resource security in Windows technologies—the assigning of access permissions to users or groups. Microsoft Windows Server 2003 offers enhancements, nuances, tools, and capabilities beyond the feature set of Windows 2000 and Windows XP, and strikingly different than Windows NT 4. Each of these additions will affect the best practices for managing and troubleshooting files and folders.

In this chapter, you will review the concepts and skills related to managing shared folders, and examine the useful Shared Folders snap-in. You will explore the Access Control List Editor, or ACL editor, with its multiple dialog boxes, each of which supports important functionality. After examining a variety of permission

configurations, you will evaluate *effective permissions*, the resulting set of permissions for a user based on user and group permissions, you will configure auditing to monitor for specific file access and operations. Finally, you will turn to IIS, which, like the File and Print Sharing service, offers another way to provide network access to files and folders.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to computer accounts in the Microsoft Active Directory directory service.. If you want hands-on practice, using the examples and lab exercises in the chapter, prepare the following:

- A Windows Server 2003 (Standard or Enterprise Edition) installed as Server01 and configured as a domain controller in the *contoso.com* domain.
- First-level organizational units (OUs): Security Groups and Employees
- The Domain Users group must be a member of Print Operators so that, during lab exercises, "normal" users can log on to a domain controller.
- Five domain local security groups in the Security Groups OU: Project 101 Team, Project 102 Team, Engineers, Managers, and Project Contractors.
- User accounts in the Employees OU for Scott Bishop, Danielle Tiedt, and Lorrin Smith-Bates, with Scott Bishop belonging to the Engineers, Project Contractors and Project 101 Team groups, Danielle Tiedt belonging to the Engineers and Project 101 Team, and Lorrin Smith-Bates belonging to the Managers and Project 101 Team.
- Access to the Shared Folders snap-in through the Computer Management console, File Server Management console (available via Manage Your Server), or a custom MMC console.

Lesson 1: Setting Up Shared Folders

We would not have networks, or our jobs, if organizations did not find it valuable to provide access to information and resources stored on one computer to users of another computer. Creating a shared folder to provide such access is therefore among the most fundamental tasks for any network administrator. Windows Server 2003 shared folders are managed with the Shared Folders snap-in.

After this lesson, you will be able to

- Create a shared folder with Windows Explorer and the Shared Folders snap-in
- Configure permissions and other properties of shared folders
- Manage user sessions and open files

Estimated lesson time: 15 minutes

Sharing a Folder

Sharing a folder configures the File And Printer Sharing For Microsoft Networks service (also known as the Server service) to allow network connections to that folder and its subfolders by clients running the Client For Microsoft Networks (also known as the Workstation service). You certainly have shared a folder using Windows Explorer by right-clicking a folder, choosing Sharing And Security, and selecting Share This Folder.

However, the familiar Sharing tab of a folder's properties dialog box in Windows Explorer is available only when you configure a share while logged on to a computer interactively or through terminal services. You cannot share a folder on a *remote* sys tem using Windows Explorer. Therefore, you will examine the creation, properties, configuration, and management of a shared folder using the Shared Folders snap-in, which can be used on both local and remote systems.

When you open the Shared Folders snap-in, either as a custom MMC console snap-in or as part of the Computer Management or File Server Management consoles, you will immediately notice that Windows Server 2003 has several default administrative shares already configured. These shares provide connection to the system directory (typically, C:\Windows) as well as to the root of each fixed hard disk drive. Each of these shares uses the dollar sign ($) in the share name. The dollar sign at the end of a share name configures the share as a *hidden share* that will not appear on browse lists, but that you may connect to with a Universal Naming Convention (UNC) in the form \\*servername**sharename*$. Only administrators can connect to the administrative shares.

To share a folder on a computer, connect to the computer using the Shared Folders snap-in by right-clicking the root Shared Folders node and choosing Connect To Another Computer. Once the snap-in is focused on the computer, click the Shares node

and, from the shortcut or Action menu, choose New Share. The important pages and settings exposed by the wizard are

- **The Folder Path page** Type the path to the folder on the *local hard drives* so, for example, if the folder is located on the server's D drive, the folder path would be D:*foldername*.
- **The Name, Description, and Settings page** Type the share name. If your network has any down-level clients (those using DOS-based systems), be sure to adhere to the 8.3 naming convention to ensure their access to the shares. The share name will, with the server name, create the UNC to the resource, in the form *servername**sharename*. Add a dollar sign to the end of the share name to make the share a hidden share. Unlike the built-in hidden administrative shares, hidden shares that are created manually can be connected to by any user, restricted only by the share permissions on the folder.
- **The Permissions page** Select the appropriate share permissions.

Managing a Shared Folder

The Shares node in the Shared Folders snap-in lists all shares on a computer and pro vides a context menu for each share that enables you to stop sharing the folder, open the share in Windows Explorer, or configure the share's properties. All the properties that you are prompted to fill out by the Share A Folder Wizard can be modified in the share's Properties dialog box, illustrated in Figure 6-1.

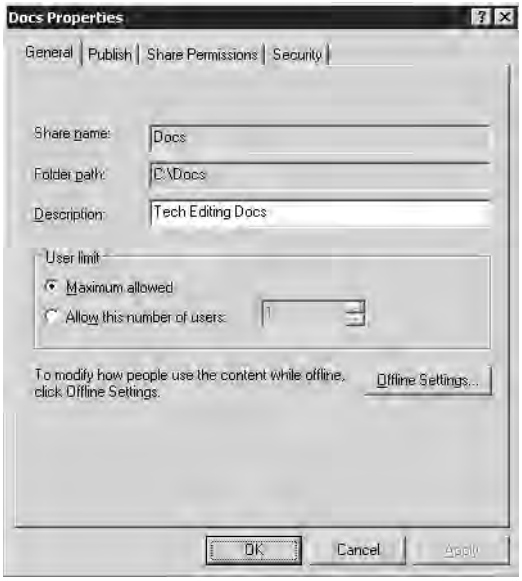

Figure 6-1 The General tab of a shared folder

The Properties tabs in the dialog box are

- **General** The first tab provides access to the share name, folder path, description, the number of concurrent user connections, and offline files settings. The share name and folder path are read-only. To rename a share, you must first stop sharing the folder then create a share with the new name.
- **Publish** If you select Publish This Share In Active Directory (as shown in Figure 6-2), an object is created in Active Directory to represent the shared folder.

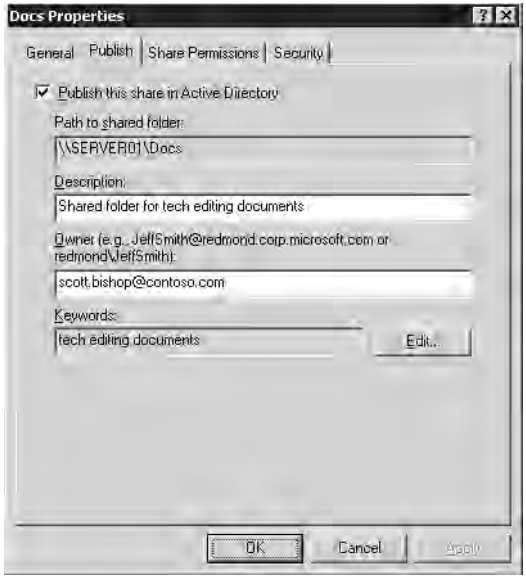

Figure 6-2 The Publish tab of a shared folder

The object's properties include a description and keywords. Administrators can then locate the shared folder based on its description or keywords, using the Find Users, Contacts and Groups dialog box. By selecting Shared Folders from the Find drop-down list, this dialog box becomes the Find Shared Folders dialog box shown in Figure 6-3.

- **Share Permissions** The Share Permissions tab allows you to configure share permissions.
- **Security** The Security tab allows you to configure NTFS permissions for the folder.

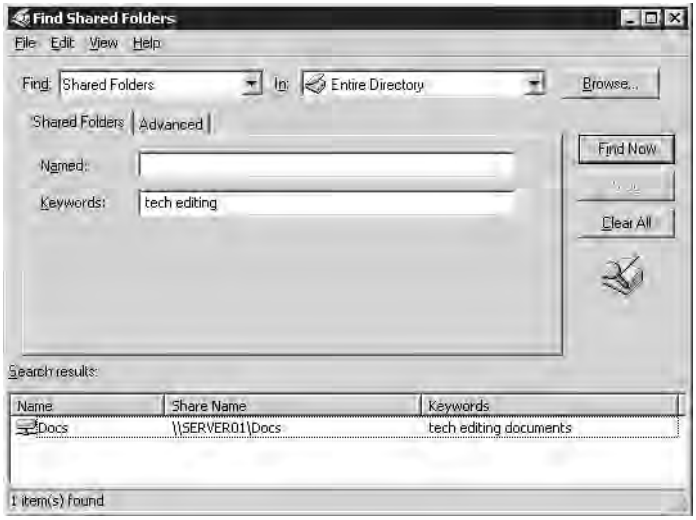

Figure 6-3 Searching for a shared folder

Configuring Share Permissions

Available share permissions are listed in Table 6-1. While share permissions are not as detailed as NTFS permissions, they allow you to configure a shared folder for funda mental access scenarios: Read, Change, and Full Control.

Table 6-1 Share Permissions

Permissions	Description
Read	Users can display folder names, file names, file data and attributes. Users can also run program files and access other folders within the shared folder.
Change	Users can create folders, add files to folders, change data in files, append data to files, change file attributes, delete folders and files, and perform actions permitted by the Read permission.
Full Control	Users can change file permissions, take ownership of files, and perform all tasks allowed by the Change permission.

Share permissions can be allowed or denied. The effective set of share permissions is the cumulative result of the Allow permissions granted to a user and all groups to which that user belongs. If, for example, you are a member of a group that has Read permission and a member of another group that has Change permission, your effective permissions are Change. However, a Deny permission will override an Allow permis sion. If, on the other hand, you are in one group that has been allowed Read access and in another group that has been denied Full Control, you will be unable to read the files or folders in that share.

Share permissions define the *maximum* effective permissions for all files and folders beneath the shared folder. Permissions can be further restricted, but cannot be broadened, by NTFS permissions on specific files and folders. Said another way, a user's access to a file or folder is the most restrictive set of effective permissions between share permissions and NTFS permissions on that resource. If you want a group to have full control of a folder and have granted full control through NTFS permissions, but the share permission is the default (Everyone: Allow Read) or even if the share permission allows Change, that group's NTFS full control access will be limited by the share permission. This dynamic means that share permissions add a layer of complexity to the management of resource access, and is one of several reasons that organizations cite for their directives to configure shares with open share permissions (Everyone: Allow Full Control), and to use only NTFS permissions to secure folders and files. See the "Three Views of Share Permissions" sidebar for more information about the variety of perspectives and drivers behind discussions of share permissions.

Three Views of Share Permissions

It is important to understand the perspectives from which share permissions are addressed in real-world implementations by Microsoft and by certification objectives and resources such as this book.

Share Permission Limitations

Share permissions have significant limitations, including the following:

- **Scope** Share permissions apply only to network access through the Client for Microsoft Networks; they do not apply to local or terminal service access to files and folders, nor to other types of network access, such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), Telnet, and so on.
- **Replication** Share permissions do not replicate through file replication service (FRS).
- **Resiliency** Share permissions are not included in a backup or restore of a data volume.
- **Fragility** Share permissions are lost if you move or rename the folder that is shared.
- **Lack of detailed control** Share permissions are not granular; they provide a single permissions template that applies to every file and folder beneath the shared folder. You cannot enlarge access to any folder or file beneath the shared folder; and you cannot further restrict access without turning to NTFS permissions.
- **Auditing** You cannot configure auditing based on share permissions.

- **The grass is truly greener** We have NTFS permissions, which are designed to provide solid, secure access control to files and folders. NTFS permissions do replicate, are included in a backup and restore of a data volume, can be audited, and provide extraordinary flexibility as well as ease of management. So organizations rely on NTFS permissions for resource access control.
- **Complexity** If both share permissions and NTFS permissions are applied, the most restrictive permission set will be effective, adding a layer of com plexity to analyzing effective permissions and troubleshooting file access.

Real-World Use of Share Permissions

Because of these limitations, the use of share permissions does not occur except for the extraordinarily rare case in which a drive volume is FAT or FAT32, which then does not support NTFS permissions. Otherwise, the "real-world" rule is: Configure shares with Everyone: Allow Full Control share permissions, and lock down the shared folder, and any other files or folders beneath it, using NTFS permissions.

Microsoft's Tightening of Share Permissions

Before Windows XP, the default share permission was Everyone: Allow Full Con trol. Using such a default, adhering to "real-world" policies was simple: adminis trators didn't change the share permission, but went straight to configuring NTFS permissions. Windows Server 2003 sets Everyone: Allow Read and Administra tors: Allow Full Control as the default share permission. This is problematic because, for all non-administrators, the entire shared folder tree is now restricted to read access.

Microsoft made this change with a noble goal: to increase security by restricting the extent to which resources are vulnerable by default when they are shared. Many administrators have shared a folder then forgotten to check NTFS permis sions only to discover, too late, that a permission was too "open." By configuring the share with read permission, Microsoft helps administrators avoid this prob lem. Unfortunately, most organizations avoid share permissions, due to their lim itations, and focus instead on providing security through NTFS permissions. Now administrators must remember to configure share permissions (to allow Everyone Full Control) to return to best practices laid out by their organizations.

Certification Objectives

There is a third perspective on share permissions: certification objectives. Although share permissions are typically implemented in accordance with strict enterprise policies (Everyone is allowed Full Control), the fact that share permis sions might one day deviate from that setting, and the possibility that data might be stored on a FAT or FAT32 volume, for which share permissions are the only

viable option for access control, means that you must understand share permissions to meet the objectives of the MCSA and MCSE exams. Of particular importance are scenarios in which both share permissions and NTFS permissions are applied to a resource, in which case the most restrictive effective permission set becomes the effective permissions set for the resource when it is accessed by a Client For Microsoft Networks service.

So pay attention to share permissions. Learn their nuances. Know how to evaluate effective permissions in combination with NTFS permissions. Then configure your shares according to your organization's guidelines, which will most likely be, unlike the new default share permission in Windows Server 2003, to allow Everyone Full Control.

Managing User Sessions and Open Files

Occasionally, a server must be taken offline for maintenance, backups must be run, or other tasks must be performed that require users to be disconnected and any open files to be closed and unlocked. Each of these scenarios will use the Shared Folders snap-in.

The Sessions node of the Shared Folders snap-in allows you to monitor the number of users connected to a particular server and, if necessary, to disconnect the user. The Open Files node enumerates a list of all open files and file locks for a single server, and allows you to close one open file or disconnect all open files.

Before you perform any of these actions, it is useful to notify the user that the user will be disconnected, so that the user has time to save any unsaved data. You can send a console message by right-clicking the Shares node. Messages are sent by the Messenger Service using the computer name, not the user name. The default state of the Messenger service in Windows Server 2003 is disabled. The Messenger service must be configured for Automatic or Manual startup and must be running before a computer can send console messages.

Practice: Setting Up Shared Folders

In this practice, you will configure a shared folder and modify the share permissions. You will then connect to the share and simulate the common procedures used before taking a server offline.

Exercise 1: Share a Folder

1. Create a folder on your C drive called Docs. Do *not* share the folder yet.
2. Open the Manage Your Server page from Administrative Tools.

3. In the File Server category, click Manage This File Server. If your server is not configured with the File Server role, you can add the role or launch the File Server Management console using the following Tip.

Tip The File Server Management console is a really nice console, so you might want to create a shortcut to it for easier access. The path to the console is %*SystemRoot*%\System32\Filesvr.msc.

4. Select the Shares node.
5. Choose Add A Shared Folder from the task list in the details pane. There are equivalent commands for adding a shared folder in the Action and the shortcut menus as well.
6. The Share A Folder Wizard appears. Click Next.
7. Type the path **c:\docs** and then click Next.
8. Accept the default share name, docs, and then click Next.
9. On the Permissions page, click Use Custom Share And Folder Permissions and then click Customize.
10. Click the check box to Allow Full Control and then click OK.
11. Click Finish, and then click Close.

Exercise 2: Connect to a Shared Folder

1. In the File Server Management console, click the Sessions node. If the node shows any sessions, click Disconnect All Sessions, from the task list, and then click Yes to confirm.
2. Choose the Run command from the Start menu. Type the UNC to the shared folder **\\server01\docs**, and then click OK.

   By using a UNC rather than a physical path, such as c:\docs, you create a network connection to the shared folder, just as a user would.
3. In the File Server Management console, click the Sessions node. Notice you are now listed as maintaining a session with the server. You may need to refresh the console by pressing F5 to see the change.
4. Click the Open Files node. Notice that you are listed as having c:\docs open.

Exercise 3: Simulate Preparing to Take a Server Offline

1. Right-click the Shares node in the File Server Management console and, from the All Tasks menu, choose Send Console Message.

Tip The Messenger service must be running on the computers that are to receive the message. Because it is not expected that a human being will be interactively logged on to the console of a server, the Messenger service is disabled by default. To send a message to yourself in this exercise, you must use the Services console to configure the Messenger service to start automatically or manually, and then start the service.

2. Type a message indicating that the server is being taken offline and that users should save their work.
3. Click Send.

   If you have a second system available, you can simulate the scenario more realis tically by connecting to the docs share and sending a message to that system.
4. Click the Open Files node.
5. Select the c:\docs file that is opened through your connection to the shared folder.
6. Close the open file. There are appropriate commands in the Action menu, the task list, and the shortcut menu.
7. Select the Sessions node.
8. Click Disconnect All Sessions in the task list. At this point, you can take the file server offline.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Which of the following tools allows you to administer a share on a remote server? Select all that apply.
   a. The Shared Folders snap-in.
   b. Windows Explorer running on the local machine, connected to the remote server's share or hidden drive share.
   c. Windows Explorer running on the remote machine in a Terminal Services or Remote Desktop session.
   d. The File Server Management console.

2. A folder is shared on a FAT32 volume. The Project Managers group is given Allow Full Control permission. The Project Engineers group is given Allow Read permis sion. Julie belongs to the Project Engineers group. She is promoted and is added to the Project Managers group. What are her effective permissions to the folder?

3. A folder is shared on a NTFS volume, with the default share permissions. The Project Managers group is given Allow Full Control NTFS permission. Julie, who belongs to the Project Managers group, calls to report problems creating files in the folder. Why can't Julie create files?

Lesson Summary

- Windows Explorer can only be used to configure shares on a local volume. This means you must be logged on locally (interactively) to the server, or using Remote Desktop (terminal services) to use Explorer to manage shares.
- The Shared Folders snap-in allows you to manage shares on a local or remote computer.
- You can create a hidden share that does not appear on browse lists by adding a dollar sign ($) to the end of the share name. Connections to the share use the UNC format: \\*servername**sharename*$.
- Share permissions define the maximum effective permissions for all files and fold ers accessed by the Client for Microsoft Networks connection to the shared folder.
- Share permissions do not apply to local (interactive), terminal services, IIS, or other types of access.

Lesson 2: Configuring File System Permissions

Windows servers support granular or detailed control of access to files and folders through NTFS. Resource access permissions are stored as access control entries (ACEs) on an ACL that is part of the security descriptor of each resource. When a user attempts to access a resource, the user's security access token, which contains the security identifiers (SIDs) of the user's account and group accounts, is compared to the SIDs in the ACEs of the ACL. This process of *authorization* has not changed fundamentally since Windows NT was introduced. However, the details of the implementation of authorization, the tools available to manage resource access, and the specificity with which you can configure access have changed with each release of Windows.

This lesson will explore the nuances and new features of Windows Server 2003's resource access control. You will learn how to use the ACL editor to manage permissions templates, inheritance, special permissions, and how to evaluate resulting effective permissions for a user or group.

After this lesson, you will be able to

- Configure permissions with the Windows Server 2003 ACL editor
- Manage ACL inheritance
- Evaluate resulting, or effective permissions
- Verify effective permissions
- Change ownership of files and folders
- Transfer ownership of files and folders

Estimated lesson time: 30 minutes

Configuring Permissions

Windows Explorer is the most common tool used to initiate management of resource access permissions, both on a local volume as well as on a remote server. Unlike shared folders, Windows Explorer can configure permissions locally and remotely.

The Access Control List Editor

As in earlier versions of Windows, security can be configured for files and folders on any NTFS volume by right-clicking the resource and choosing Properties (or Sharing And Security) then clicking the Security tab. The interface that appears has many aliases; it has been called the Permissions dialog box, the Security Settings dialog box, the Security tab or the Access Control List editor (ACL editor). Whatever you call it, it looks the same. An example can be seen in the Security tab of the Docs Properties dialog box, as shown in Figure 6-4.

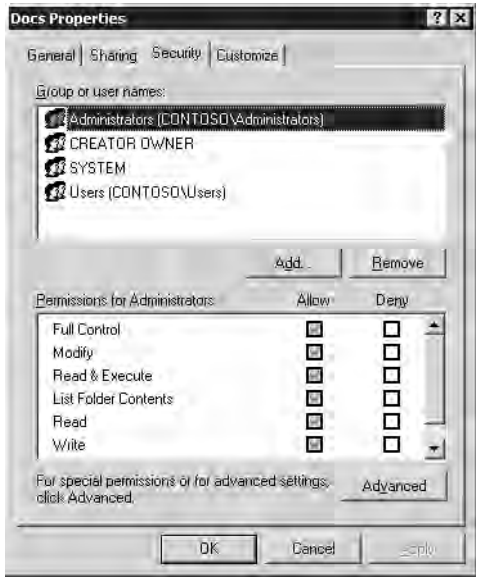

Figure 6-4 The ACL editor in the Docs Properties dialog box

Prior to Windows 2000, permissions were fairly simplistic, but with Windows 2000 and later versions, Microsoft enabled significantly more flexible and powerful control over resource access. With more power came more complexity, and now the ACL editor has three dialog boxes, each of which supports different and important functionality.

The first dialog box provides a "big picture" view of the resource's security settings or permissions, allowing you to select each account that has access defined and to see the permissions templates assigned to that user, group, or computer. Each template shown in this dialog box represents a bundle of permissions that together allow a commonly configured level of access. For example, to allow a user to read a file, several granular permissions are needed. To mask that complexity, you can simply apply the Allow:Read & Execute permissions template and, behind the scenes, Windows sets the correct file or folder permissions.

To view more details about the ACL, click Advanced, which exposes the second of the ACL editor's dialog boxes, the Advanced Security Settings For Docs dialog box, as shown in Figure 6-5. This dialog box lists the specific access control entries that have been assigned to the file or folder. The listing is the closest approximation in the user interface to the actual information stored in the ACL itself. The second dialog also enables you to configure auditing, manage ownership, and evaluate effective permissions.

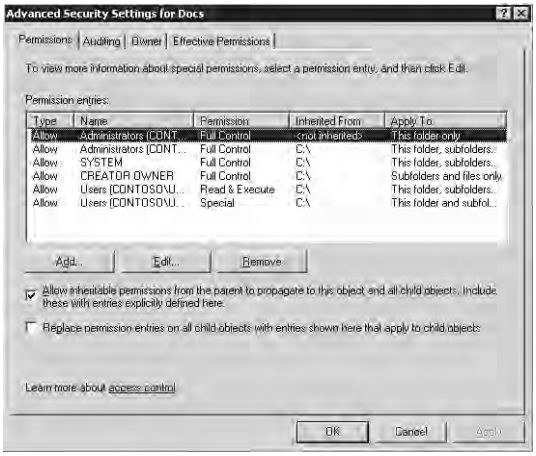

Figure 6-5 The ACL editor's Advanced Security Settings dialog box

If you select a permission in the Permission Entries list and click Edit, the ACL editor's third dialog box appears. This Permission Entry For Docs dialog box, shown in Figure 6-6, lists the detailed, most granular permissions that comprise the permissions entry in the second dialog box's Permissions Entries list and the first dialog box's Permissions For Users list.

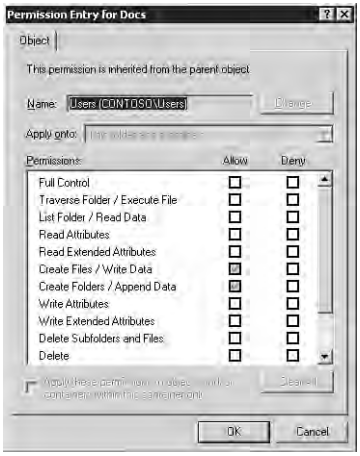

Figure 6-6 The ACL editor's Permission Entry dialog box

Exam Tip The Shared Folders snap-in also allows you to access the ACL editor. Open the properties of a shared folder and click the Security tab.

Adding and Removing Permission Entries

Any security principal may be granted or denied resource access permissions. In Windows Server 2003, the valid security principals are: users, groups, computers, and the special InetOrgPerson object class (described in RFC 2798), which is used to represent users in certain cross-directory platform situations. To add a permission, click the Add button on either the first or second ACL editor dialog box. The Select User, Computer Or Group dialog box will help you identify the appropriate security principal. Then select appropriate permissions. The interface has changed slightly from previous versions of Windows, but not enough to prevent an experienced administrator from mastering the new user interface quickly. You can remove an explicit permission that you have added to an ACL by selecting the permission and clicking Remove.

Modifying Permissions

A permission may be modified in the dialog box by selecting or clearing the Allow or Deny check boxes on the Security tab to apply permissions templates.

For a finer degree of control, click Advanced, select a permission entry and click Edit. Only explicit permissions may be edited. Inherited permissions are discussed later in this lesson.

The Permission Entry For Docs dialog box, shown in Figure 6-6, will allow you to mod ify permissions and specify the scope of the permissions inheritance, through the Apply Onto drop-down list.

Caution Be certain that you understand the impact of changes you make in this dialog box. You can be grateful for the detailed control Microsoft has enabled, but with increased granu larity comes increased complexity and increased potential for human error.

New Security Principals

Windows Server 2003, unlike Windows NT 4, allows you to add *computers* or groups of computers to an ACL, thereby adding flexibility to control resource access based on the client computer, regardless of the user who attempts access. For example, you may want to provide a public computer in the employee lounge, but prevent a manager from exposing sensitive data during his or her lunch break. By adding the computer to ACLs and denying access permission, the manager who can access sensitive data from his or her desktop is prevented from accessing it from the lounge.

Windows Server 2003 also allows you to manage resource access based on the type of logon. You can add the special accounts, Interactive, Network, and Terminal Server User to an ACL. Interactive represents any user logged on locally to the console. Ter minal Server User includes any user connected via remote desktop or terminal services.

Network represents a connection from the network, for example a Windows system running Client for Microsoft Networks.

Permissions Templates and Special Permissions

Permissions templates, visible on the Security tab in the first dialog box are bundles of special permissions, which are fully enumerated in the third dialog box, Permissions Entry For Docs. Most of the templates and special permissions are self-explanatory, while others are beyond the scope of this book. However, the following points are worth noting:

- **Read & Execute** This permissions template is sufficient to allow users to open and read files and folders. Read & Execute will also allow a user to copy a resource, assuming they have permission to write to a target folder or media. There is no permission in Windows to prevent copying. Such functionality will be possible with Digital Rights Management technologies as they are incorporated into Windows platforms.
- **Write and Modify** The Write permissions template applied to a folder allows users to create a new file or folder (when applied to a folder) and, when applied to a file, to modify the contents of a file as well as its attributes (hidden, system, read-only) and extended attributes (defined by the application responsible for the document). The Modify template adds the permission to delete the object.
- **Change Permissions** After modifying ACLs for a while, you might wonder who can modify permissions. The answer is, first, the owner of the resource. Ownership will be discussed later in this lesson. Second, any user who has an effective permission that allows Change Permission can modify the ACL on the resource. The Change Permission must be managed using the ACL editor's third dialog box, Permission Entry For Docs. It is also included in the Full Control permission template.

Inheritance

Windows Server 2003 supports permissions inheritance, which simply means that permissions applied to a folder will, by default, apply to the files and folders beneath that folder. Any change to the parent's ACL will similarly affect all contents of that folder. Inheritance enables you to create single points of administration, managing a single ACL on a branch or resources under a folder.

Understanding Inheritance

Inheritance is the result of two characteristics of a resource's security descriptor. First, permissions are, by default, inheritable. As previously shown in Figure 6-5, the permis sion Allow Users to Read & Execute is specified to Apply to: This folder, subfolders,

and files. That alone, however, is not enough to make inheritance work. The other half of the story is that new objects, when created, are set by default to "Allow Inheritable Permissions From The Parent To Propagate To This Object..." the check box visible in the same figure.

So a newly created file or folder will inherit the inheritable permissions from its parent, and any changes to the parent will affect the child files and folders as well. It is helpful to understand this two-step implementation of inheritance because it gives us two ways to manage inheritance: from the parent and from the child.

Inherited permissions are displayed differently in each dialog box of the ACL editor. The first and third dialog boxes (Security tab and Permissions Entry For Docs) show inherited permissions as dimmed check marks, to distinguish them from permissions that are set directly on the resource, called explicit permissions, which are not dimmed. The second dialog box (Advanced Security Settings) shows, for each permission entry, from what folder the permission entry is inherited.

Overriding Inheritance

Inheritance allows you to configure permissions high in a folder tree. Such initial permissions, and any changes to those permissions, will propagate to all the files and fold ers in that tree that are, by default, configured to allow inheritance.

Occasionally, however, you might need to modify permissions on a subfolder or file, to provide additional access or restrict access to a user or group. You cannot remove inherited permissions from an ACL. You can override an inherited permission by assigning an explicit permission. Alternatively, you can block all inheritance and create an entirely explicit ACL.

To override an inherited permission by assigning an explicit permission, simply check the appropriate permissions box. For example, if a folder has an inherited Allow Read permission assigned to the Sales Reps group, and you do not want Sales Reps to access the folder, you can select the box to Deny Read.

To override all inheritance, open the resources Advanced Security Settings dialog box and clear Allow Inheritable Permissions From The Parent To Propagate To This Object... You will block all inheritance from the parent. You will then have to manage access to the resource by assigning sufficient explicit permissions.

To help you create an explicit permissions ACL, Windows gives you a choice when you choose to disallow inheritance. You are asked whether you want to Copy or Remove permissions entries, as shown in Figure 6-7.

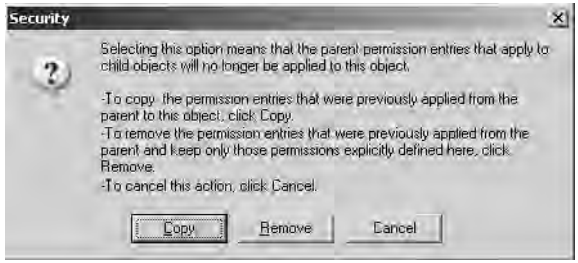

Figure 6-7 Copying or removing permissions entries

Copy will create explicit permissions identical to what was inherited. You can then remove individual permissions entries that you do not want to affect the resource. If you choose Remove, you will be presented with an empty ACL, to which you will add permissions entries. The result is the same either way; an ACL populated with explicit permissions. The question is whether it is easier to start with an empty ACL and build it from scratch or start with a copy of the inherited permissions and modify the list to the desired goal. If the new ACL is wildly different than the inherited permissions, choose Remove. If the new ACL is only slightly different than the result of inherited permissions, it is more efficient to choose Copy.

When you disallow inheritance by deselecting the Allow Inheritable Permissions option, you block inheritance. All access to the resource is managed by explicit permissions assigned to that file or folder. Any changes to the ACL of its parent folder will *not* affect the resource; although the parent permissions are inheritable, the child does not inherit. Block inheritance sparingly because it increases the complexity of manag ing, evaluating, and troubleshooting resource access.

Reinstating Inheritance

Inheritance can be reinstated in two ways: from the child resource or from the parent folder. The results differ slightly. You might reinstate inheritance on a resource if you disallowed inheritance accidentally or if business requirements have changed. Simply re-select the Allow Inheritable Permissions option in the Advanced Security Settings dialog box. Inheritable permissions from the parent will now apply to the resource. All explicit permissions you assigned to the resource remain, however. The resulting ACL is a combination of the explicit permissions, which you might choose to remove, and the inherited permissions. Because of this dynamic, you might not see some inherited permissions in the first or third ACL editor dialog boxes. For example, if a resource has an explicit permission, Allows Sales Reps Read & Execute, and the parent folder has the same permission, when you choose to allow inheritance on the child, the result will be that the child has both an inherited *and* an explicit permission. You will see a check mark in the first and third dialog boxes; the explicit permission obscures the inherited permission in the interface. But the inherited permission is actually present, which can be confirmed in the second dialog box, Advanced Security Settings.

The second method for reinstating inheritance is from the parent folder. In the Advanced Security Settings dialog box of a folder, you may select the check box, Replace Permission Entries On All Child Objects With Entries Shown Here That Apply To Child Objects. The result: all ACLs on subfolders and files are removed. The permissions on the parent are applied. You might see this as "blasting through" the parent's permissions. After applying this option, any explicit permission that had been applied to subfolders and files is removed, unlike the method used for reinstating inheritance on the child resources. Inheritance is restored, so any changes to the parent-folder ACL are propagated to its subfolders and files. At this point, you might set new, explicit permissions on subfolders or files. The Replace Permissions option does its job when you apply it, but does not continuously enforce parent permissions.

Effective Permissions

It is common for users to belong to more than one group, and for those groups to have varying levels of resource access. When an ACL contains multiple entries, you must be able to evaluate the permissions that apply to a user based on his or her group memberships. The resulting permissions are called *effective permissions*.

Exam Tip Effective permissions are a common exam objective on most of the Microsoft Windows Server 2003 core exams, as well as on design and client exams. Pay close attention to this information, and to any practice questions regarding effective permissions so you can be certain you have mastered the topic.

Understanding Effective Permissions

The rules that determine effective permissions are as follows:

- **File permissions override folder permissions.** This isn't really a rule, but it is often presented that way in documentation, so it is worth addressing. Each resource maintains an ACL that is solely responsible for determining resource access. Although entries on that ACL may appear because they are inherited from a parent folder, they are nevertheless entries on that resource's ACL. The security subsystem does not consult the parent folder to determine access at all. So you may interpret this rule as: The only ACL that matters is the ACL on the resource.
- **Allow permissions are cumulative.** Your level of resource access may be determined by permissions assigned to one or more groups to which you belong. The Allow permissions that are assigned to any of the user, group, or computer IDs in your security access token will apply to you, so your effective permissions are fundamentally the sum of those Allow permissions. If the Sales Reps group is allowed Read & Execute and Write permissions to a folder, and the Sales Managers group is allowed Read & Execute and Delete permissions, a user who belongs to

both groups will have effective permissions equivalent to the Modify permissions template: Read & Execute, Write and Delete.

- **Deny permissions take precedence over Allow permissions.** A permission that is denied will override a permission entry that allows the same access. Extending the example above, if the Temporary Employees group is denied Read permission, and a user is a temporary sales representative, belonging to both Sales Reps and Temporary Employees, that user will not be able to read the folder.

Note Best practice dictates that you minimize the use of Deny permissions and focus instead on allowing the minimal resources permissions required to achieve the business task. Deny permissions add a layer of complexity to the administration of ACLs, and should be used only where absolutely necessary to exclude access to a user who has been granted permissions to the resource through other group memberships.

Exam Tip If a user is unable to access a resource due to a Deny permission, but access is desired, you must either remove the Deny permission or remove the user from the group to which the Deny permission is applied. If the Deny permission is inherited, you may provide access by adding an explicit Allow permission.

- **Explicit permissions take precedence over inherited permissions.** A permission entry that is explicitly defined for a resource will override a conflicting inherited permission entry. This follows common-sense design principles: A parent folder sets a "rule" through its inheritable permissions. A child object requires access that is an exception to the rule, and so an explicit permission is added to its ACL. The explicit permission takes precedence.

Exam Tip A result of this dynamic is that an *explicit Allow permission will override an inherited Deny permission*.

Evaluating Effective Permissions

Complexity is a possibility, given the extraordinary control over granular permissions and inheritance that NTFS supports. With all those permissions, users and groups, how can you know what access a user actually has?

Microsoft added a long-awaited tool to help answer that question. The Effective Permissions tab of the Advanced Security Settings dialog box, shown in Figure 6-8, provides a reliable approximation of a user's resulting resource access.

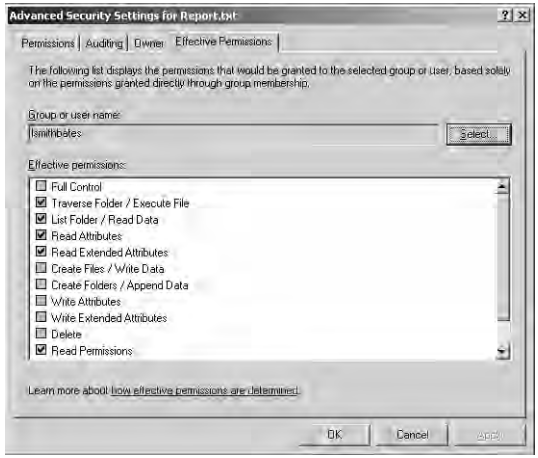

Figure 6-8 The Effective Permissions tab of the Advanced Security Settings dialog box

To use the Effective Permissions tool, click Select and identify the user, group, or built-in account to analyze. Windows Server 2003 then produces a list of effective permissions. This list is an approximation only. It does not take share permissions into account, nor does it evaluate the account's special memberships, such as the following:

- Anonymous Logon
- Batch
- Creator Group
- Dialup
- Enterprise Domain Controllers
- Interactive
- Network
- Proxy
- Restricted
- Remote Interactive Logon
- Service
- System
- Terminal Server User
- Other Organization
- This Organization

An ACL can contain entries for the Network or Interactive accounts, for example, which would provide the opportunity for a user to experience different levels of resource access depending on whether the user was logged on to the machine or using a network client. Because the user in question is not logged on, logon-specific permissions entries are ignored. However, as an extra step, you can evaluate effective permissions for a built-in or special account such as Interactive or Network.

Resource Ownership

Windows Server 2003 includes a special security principal called Creator Owner, and an entry in a resource's security descriptor that defines the object's owner. To fully manage and troubleshoot resource permissions, you must understand these two parts of the security picture.

Creator Owner

When a user creates a file or folder (which is possible if that user is allowed Create Files/Write Data or Create Folders/Append Data, respectively), the user is the creator and initial owner of that resource. Any permissions on the parent folder assigned to the special account Creator Owner are explicitly assigned to the user on the new resource.

As an example, assume that a folder allows users to create files (allow Create Files/Write Data), and the folder's permissions allows users to Read & Execute, and Creator Owner Full Control. This permission set would allow Maria to create a file. Maria, as the creator of that file, would have full control of that file. Tia can also create a file, and would have full control of her file. However, Tia and Maria would only be able to read each other's files. Tia could, however, change the ACL on the file she created. Full Control includes the Change Permission.

Ownership

If for some reason Tia managed to modify the ACL and deny herself Full Control, she could nevertheless modify the ACL, because an object's owner can always modify its ACL, preventing users from permanently locking themselves out of their files and folders.

It is best practice to manage object ownership so that an object's owner is correctly defined. This is partly because owners can modify ACLs of their objects, and also because newer technologies, such as disk quotas, rely on the ownership attribute to calculate disk space used by a particular user. Prior to Windows Server 2003, managing ownership was awkward. Windows Server 2003 has added an important tool to simplify ownership transfer.

An object's owner is defined in its security descriptor. The user who creates a file or folder is its initial owner. Another user can take ownership, or be given ownership of the object using one of the following processes:

- **Administrators can take ownership.** A user who belongs to the Administrators group of a system, or who has otherwise been granted the Take Ownership user right, can take ownership of any object on the system.

  To take ownership of a resource, click the Owner tab of the Advanced Security Settings dialog box, as shown in Figure 6-9. Select your user account from the list and click Apply. Select the Replace Owner On Subcontainers And Objects check box to take ownership of subfolders and files.

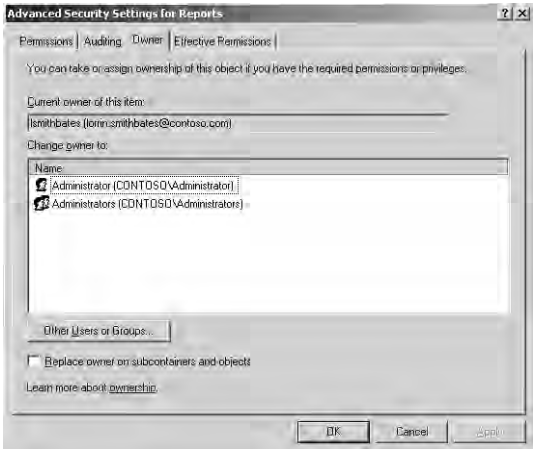

Figure 6-9 The Owner tab of the Advanced Security Settings dialog box

- **Users can take ownership if they are allowed Take Ownership permission.** The special permission Take Ownership can be granted to any user or group. A user with an Allow Take Ownership permission can take ownership of the resource and then, as owner, modify the ACL to provide sufficient permissions.
- **Administrators can facilitate the transfer of ownership.** An administrator can take ownership of any file or folder. Then, as owner, the administrator can change permissions on the resource to grant Allow Take Ownership permission to the new owner, who then can take ownership of the resource.
- **Restore Files And Directories user right enables the transfer of ownership.** A user with the Restore Files And Directories rights may transfer ownership of a file from one user to another. If you have been assigned the Restore Files And Directories right, you can click Other Users Or Groups and select the new owner. This capability is new in Windows Server 2003, and makes it possible for administrators and backup operators to manage and transfer resource ownership without requiring user intervention.

Practice: Configuring File System Permissions

In this practice, you will use the ACL editor to secure resources, evaluate effective permissions and transfer ownership of files. Be certain that you have configured the user and group accounts outlined in this chapter's "Before You Begin" section.

Exercise 1: Configuring NTFS Permissions

1. Open the c:\docs folder that was shared in Lesson 1's practice.
2. Create a folder called Project 101.
3. Open the ACL editor by right-clicking Project 101, choosing Properties, and click ing the Security tab.
4. Configure the folder so that the folder allows the access outlined in the table below. This will require you to consider and configure, inheritance and permis sions for groups.

Security Principal	Access
Administrators	Full Control
Users in the Project 101 Team	Can read data, add files and folders, and have full control of the files and folders they create.
Managers	Can read and modify all files, but cannot delete any files that they did not create. Managers should have full control of the files and folders they create.
System	Services running as the System account should have full control.

When you believe you have configured correct permissions, click Apply and click Advanced. Compare the Advanced Security Settings dialog box to the dialog box shown in Figure 6-10.

To configure these permissions, you must disallow inheritance. Otherwise, all users, not just those in the Project 101 group, will be able to read files in the Project 101 folder. The parent folder, c:\docs, is propagating the Users: Allow Read & Execute permission. The only way to prevent this access is to deselect the Allow Inheritable Permissions From The Parent… option. Notice that the requirements did not specify that you needed to prevent Users from reading, but it was also not indicated that Users required read access, and it is a security best practice to permit only the minimum required access.

After disallowing inheritance, the Advanced Security Settings dialog box should look like the dialog box in Figure 6-10.

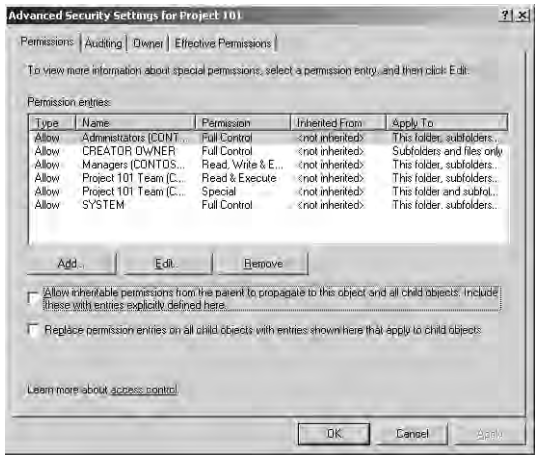

Figure 6-10 The Permissions tab of the Advanced Security Settings dialog box

The option to allow inheritance has been deselected and all permissions are shown as <not inherited>. Administrators, System, and Creator Owner have full control. Remember that when Creator Owner has full control, a user who creates a file or folder is given full control of that resource. The Project 101 group is listed as having a special permission entry. If you select that entry and click View/Edit, you will see the specific permissions assigned to the Project 101 group should match the dialog box shown in Figure 6-11.

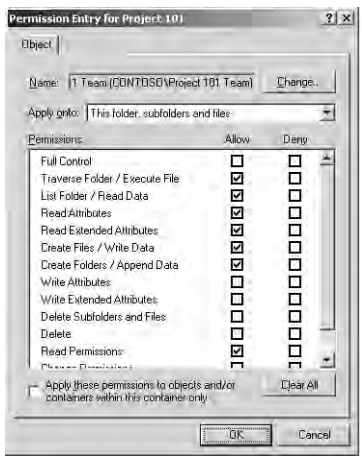

Figure 6-11 Special permissions for the Project 101 group

The Managers have Allow: Read, Write & Execute permission. This template includes the permissions to create files and folders and, like Project 101 team members, if a manager creates a resource, Managers are given the Creator Owner permissions for that resource. This permission set does *not* allow Managers to delete other users' files. Remember that the Modify permissions template, which you did not assign, *does* include the Delete permission.

Exercise 2: Working with Deny Permissions

1. Assume a group of contractors is hired. All user accounts for contractors are mem bers of the Project Contractors group, and do not belong to any other group in the domain. What must you do to prevent contractors from accessing the Project 101 folder you secured in the previous exercise?

   Nothing. Because contractors do not belong to other groups in the domain, they do not have permissions given to them by the current ACL that would allow any resource access. It is therefore not necessary to deny permissions.

2. Assume that some user accounts, such as Scott Bishop's account, belong to both the Project Contractors and the Engineers groups. What must be done to prevent access by contractors?

   In this case, you must assign Deny permissions to the Project Contractors group. Because they will receive Allow permissions assigned to other groups, you must override those permissions with Deny permissions.

3. Configure the folder to Deny Project Contractors Full Control.

Exercise 3: Effective Permissions

1. Open the Advanced Security Settings dialog box for the Project 101 folder by opening the folder's properties, clicking Security, then clicking Advanced.
2. Click Effective Permissions.
3. Select each of the following users and verify their permissions.

User	Effective Permissions
Scott Bishop	No permissions
Danielle Tiedt	Traverse Folder / Execute File List Folder / Read Data Read Attributes Read Extended Attributes Create Files / Write Data Create Folders / Append Data Read Permissions

(Continued)

User	Effective Permissions
Lorrin Smith-Bates	Traverse Folder / Execute File List Folder / Read Data Read Attributes Read Extended Attributes Create Files / Write Data Create Folders / Append Data Write Attributes Write Extended Attributes Read Permissions

If these permissions do not match yours, there is either an error in the permission list (in which case, go back to Exercises 1 and 2) or in groups and group membership (in which case, see this chapter's "Before You Begin" section). Correct any errors and reverify effective permissions until they match these.

Exercise 4: Ownership

1. Log on as Danielle Tiedt.
2. Open the shared folder by connecting to \\Server01\Docs.
3. Open the Project 101 folder and create a text file called Report.
4. Open the Advanced Security Settings dialog box for Report.
5. Confirm that all permissions are inherited from the parent folder. What differences are there in the ACL between this object and the Project 101 folder?

   The Project 101 folder grants Full Control to Creator Owner. The Report file grants Full Control to Danielle. When she created the file, her SID was assigned the permissions granted to the special Creator Owner group. In addition, the Project 101 Team's permission to Create Files and Create Folders is a folder permission, so it does not appear on the ACL of Report.

6. Log on as Administrator.
7. Open the Advanced Security Settings dialog box for Report.
8. Click Owner.
9. Confirm that Danielle is listed as the current owner.
10. Select your user account and click Apply. You are now the owner of the object.
11. A user with the Restore Files And Directories user right is able to transfer ownership to another user. Click Other Users Or Group and select Lorrin Smith-Bates. Once Lorrin's account is displayed in the Change Owner To list, select it and click Apply.
12. Confirm that Lorrin is now the owner of the Report.

13. Do you think that Lorrin now has full control of the object? Why or why not? Do you think that Danielle will keep full control, or will her permissions change? Confirm using the Effective Permissions page.

    Lorrin does not have full control—only Modify permission. Lorrin is a member of the Managers group, which has Modify permission. The Full Control permission assigned to Creator Owner is only applied to a user when the user creates an object.

Note Once an object has been created, changing ownership does not modify the ACL in any way. However, the new owner (or any user with Allow Change Permissions) can modify the ACL, as an additional step, to provide himself or herself with sufficient resource access.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What are the minimum NTFS permissions required to allow users to open documents and run programs stored in a shared folder?
   - **a.** Full Control
   - **b.** Modify
   - **c.** Write
   - **d.** Read & Execute
   - **e.** List Folder Contents
2. Bill complains that he is unable to access the department plan. You open the Security tab for the plan and you find that all permissions on the document are inherited from the plan's parent folder. There is a Deny Read permission assigned to a group to which Bill belongs. Which of the following methods would enable Bill to access the plan?
   - **a.** Modify the permissions on the parent folder by adding the permission Bill:Allow Full Control.
   - **b.** Modify the permissions on the parent folder by adding the permission Bill:Allow Read.
   - **c.** Modify the permissions on the plan by adding the permission: Bill:Allow Read.
   - **d.** Modify the permissions on the plan by deselecting Allow Inheritable Permissions, choosing Copy, and removing the Deny permission.

e. Modify the permissions on the plan by deselecting Allow Inheritable Permissions, choosing Copy, and adding the permission Bill:Allow Full Control.

f. Remove Bill from the group that is assigned the Deny permission.

3. Bill calls again to indicate that he still cannot access the departmental plan. You use the Effective Permissions tool, select Bill's account, and the tool indicates that Bill is, in fact, allowed sufficient permissions. What might explain the discrepancy between the results of the Effective Permissions tool and the issue Bill is reporting?

Lesson Summary

- NTFS permissions can be configured using the ACL editor, which itself has three dialog boxes: the Security tab, Advanced Security Settings, and Permission Entry For.
- Permissions can be allowed or denied; explicit or inherited. A Deny permission takes precedence over an Allow permission; and an explicit permission takes precedence over an inherited permission. The result is that an explicit Allow permission can override an inherited Deny permission.
- Inheritance allows an administrator to manage permissions from a single parent folder that contains files and folders that share common resource access requirements. A new object's ACL will, by default, include the inheritable permissions from the parent folder.
- It is possible to change the effect of inherited permissions on an object several ways. You can modify the original (parent's) permission and allow the new permission to be inherited by the object; you can set an explicit permission on the object, which will take precedence over the inherited permission; or you can disallow inheritance on the object and configure an ACL with explicit permissions that define resource access.
- The Effective Permissions tab of the Advanced Security Settings dialog box is a useful tool that provides an approximation of resource access for a user or a group by analyzing that account's permissions as well as the permissions of groups to which that account belongs.
- The owner of an object can modify the object's ACL at any time. A user that is allowed Take Ownership permission may take ownership of the object, and administrators may take ownership of any object on the system. Administrators, Backup Operators, and other accounts that have been given the Restore Files And Directories user right can transfer ownership of a file or folder from the current owner to any other user or group.

Lesson 3: Auditing File System Access

Many organizations elect to audit file system access to provide insight into resource uti lization and potential security vulnerabilities. Windows Server 2003 supports granular auditing based on user or group accounts and the specific actions performed by those accounts. To configure auditing, you must complete three steps: specify auditing set tings, enable audit policy, and evaluate events in the security log. This lesson will explore these three processes and provide guidance to effective auditing, so that you can leverage auditing to meet business requirements without being drowned in logged events.

After this lesson, you will be able to

- Configure audit settings on a file or folder
- Enable auditing on a standalone server or for a collection of servers
- Examine audited events in the Security log

Estimated lesson time: 20 minutes

Configuring Audit Settings

To specify the actions you wish to monitor and track, you must configure audit settings in the file's or folder's Advanced Security Settings dialog box. The Auditing tab, shown in Figure 6-12, looks strikingly similar to the Permissions tab before it. Instead of add ing permissions entries, however, you add auditing entries.

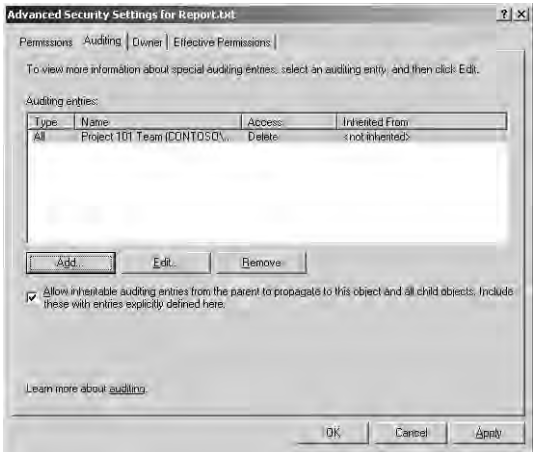

Figure 6-12 Auditing tab of the Advanced Security Settings dialog box

Click Add to select the user, group, or computer to audit. Then, in the Auditing Entry dialog box, as shown in Figure 6-13, indicate the permission uses to audit.

Figure 6-13 Auditing Entry dialog box

You are able to audit for successes, failures, or both as the account attempts to access the resource using each of the granular permissions assigned to the object.

Successes can be used to audit the following:

- To log resource access for reporting and billing.
- To monitor for access that would indicate that users are performing actions greater than what you had planned, indicating permissions are too generous.
- To identify access that is out of character for a particular account, which might be a sign that a user account has been breached by a hacker.

Auditing for failed access allows you:

- To monitor for malicious attempts to access a resource to which access has been denied.
- To identify failed attempts to access a file or folder to which a user does require access. This would indicate that permissions are not sufficient to achieve a business task.

Audit settings, like permissions, follow rules of inheritance. Inheritable auditing settings are applied to objects that allow inheritance.

Note Audit logs have the tendency to get quite large, quite rapidly, so a golden rule for auditing is to configure the bare minimum required to achieve the business task. Specifying to audit successes and failures on an active data folder for the Everyone group using Full Control (all permissions) would generate enormous audit logs that could affect the performance of the server and would make locating a specific audited event all but impossible.

Enabling Auditing

Configuring auditing entries in the security descriptor of a file or folder does not, in itself, enable auditing. Auditing must be enabled through policy. Once auditing is enabled, the security subsystem begins to pay attention to the audit settings, and to log access as directed by those settings.

Audit policy may be enabled on a stand-alone server using the Local Security Policy console, and on a domain controller using the Domain Controller Security Policy console. Select the Audit Policy node under the Local Policies node and double-click the policy, Audit Object Access. Select Define These Policy Settings and then select whether to enable auditing for successes, failures, or both.

Note Remember that the access that is audited and logged is the combination of the audit entries on specific files and folders, and the settings in Audit Policy. If you have configured audit entries to log failures, but the policy enables only logging for successes, your audit logs will remain empty.

You may also enable auditing for one or more computers using Active Directory Group Policy Objects (GPOs). The Audit Policy node is located under Computer Configuration, Windows Settings, Security Settings, Local Policies, Audit Policy. Like all group policies, the computers that are affected by the policy will be those contained within the scope of the policy. If you link a policy to the Servers OU and enable auditing, all computers objects in the Servers OU will begin to audit resource access according to audit entries on files and folders on those systems.

Examining the Security Log

Once audit entries have been configured on files or folders, and auditing object access has been enabled through local or group policy, the system will begin to log access according to the audit entries. You can view and examine the results using Event Viewer and selecting the Security log, as shown in Figure 6-14.

As you can see, the Security log can be quite busy, depending on the types of auditing being performed on the machine. You can sort the events to help you identify object access events by clicking the Category column header and locating the Object Access events.

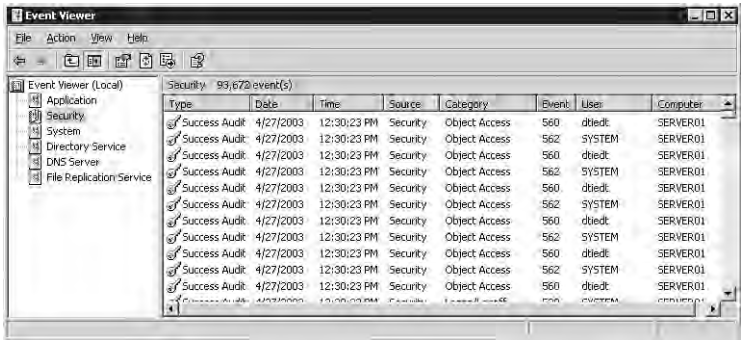

Figure 6-14 The Security log in Event Viewer

Sorting will, however, provide little assistance as you dig through the logged events. You will often be better served by filtering the event log, which can be done by choos ing the Filter command from the View menu, or alternatively by selecting the Security log, then Properties from the Action or shortcut menus, and then clicking the Filter tab. The Filter tab enables you to specify criteria including the event type, category, source, date range, user, and computer. Figure 6-15 illustrates an example of a filter applied to identify object access audit events on a specific date.

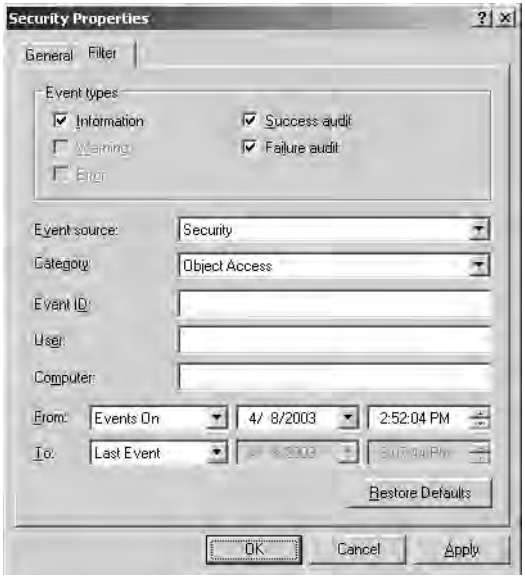

Figure 6-15 The Filter tab

Finally, you have the option to export the Security log by selecting the Save Log File As command from the log's context menu. The native event log file format takes a .evt extension. You can open that file with Event Viewer on another system. Alternatively, you can save the log to tab- or comma-delimited file formats, which can be read by a

number of analysis tools including Microsoft Excel. In Excel, you can of course apply filters as well to search for more specific information, such as the contents of the event's Description field.

Practice: Auditing File System Access

In this practice, you will configure auditing settings, enable audit policies for object access, and filter for specific events in the security log. The business objective is to monitor the deletion of files from an important folder, to ensure that only appropriate users are deleting files.

Exercise 1: Configure Audit Settings

1. Log on as Administrator.
2. Open the Advanced Security Settings dialog box for the C:\Docs\Project 101 folder.
3. Click the Auditing tab.
4. Add an audit entry to track the Project 101 Team group. Specify that you wish to monitor Success and Failure of the Delete permission.

Exercise 2: Enable Audit Policy

Because you are logged on to a domain controller, you will use the Domain Controller Security Policy console to enable auditing. On a stand-alone server you would use Local Security Policy. You could also leverage GPOs to enable auditing.

1. Open Domain Controller Security Policy from the Administrative Tools folder.
2. Expand Local Policies and select Audit Policy.
3. Double-click Audit Object Access.
4. Select Define These Policy Settings.
5. Specify to enable auditing for both success and failure audit entries.
6. Click OK, and then close the console.
7. To refresh the policy, and to ensure that all settings have been applied, open a command prompt and type the command **gpupdate**.

Exercise 3: Generate Audit Events

1. Log on as Danielle Tiedt.
2. Connect to \\Server01\Docs\Project 101.
3. Delete the Report text file.

Exercise 4: Examine the Security Log

1. Log on as Administrator.
2. Open Event Viewer from the Administrative Tools folder.
3. Select the Security log.
4. What types of events do you see in the Security log? Only Object Access events? Other types of events? Remember that policies can enable auditing for numerous security-related actions, including directory service access, account management, logon, and more.
5. To filter the log and narrow the scope of your search, choose the Filter command from the View menu.
6. Configure the filter to be as narrow as possible. What do you know about the event you are trying to locate? You know it is a success or failure audit; that it is an Object Access event category; and that it occurred today. Check your work by referring to Figure 6-15.
7. Click Apply.
8. Can you more easily locate the event that marked Danielle's deletion of the Report file? Open the event and look at its contents. The description indicates the user and the file and the action. You could not filter for contents of the description in Event Viewer, but you could do so by exporting the file to a log analysis tool or to Microsoft Excel.
9. (Optional) If you have access to Microsoft Excel, right-click the Security log node and choose Save Log File As. Enter a name and select Comma-Delimited as the file type. Open the file in Excel.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Which of the following must be done to generate a log of resource access for a file or folder? Select all that apply.
   a. Configure NTFS permissions to allow the System account to audit resource access.
   b. Configure audit entries to specify the types of access to audit.
   c. Enable the Audit Privilege Use policy.
   d. Enable the Audit Object Access policy.

2. Which of the following are valid criteria for a security log filter to identify specific file and folder access events? Select all that apply.
   a. The date of the event
   b. The user that generated the event
   c. The type of object access that generated the event
   d. Success or failure audit
3. Users at Contoso Ltd. use Microsoft Office applications to access resources on Server01. Your job is to monitor Server01 to ensure that permissions are not too restrictive, so that users are not prevented from achieving their assignments. Which log, and which type of event, will provide the information you require?
   a. Application log; Success Event
   b. Application log; Failure Event
   c. Security log; Success Event
   d. Security log; Failure Event
   e. System log; Success Event
   f. System log; Failure Event

Lesson Summary

- Audit entries are contained in the security descriptor of files and folders on NTFS volumes. They are configured using Windows Explorer, from the properties of a file or folder, using the Advanced Security Settings dialog box.
- Audit entries alone do not generate audit logs. You must also enable the Audit Object Access policy from Local Security Policy, the Domain Controller Security Policy, or a GPO.
- The Security log, viewable with the Event Viewer snap-in, allows you to locate and examine object access events.

Lesson 4: Administering Internet Information Services

Lesson 1 discussed the issues related to sharing a folder so that users, with the Client For Microsoft Networks, can access resources on a server running the File And Print Sharing For Microsoft Networks service. That is, however, only one means by which users can access the files and folders they require. It is also possible to enable access through Internet technologies such as FTP and Web (HTTP) services.

In this lesson, you will learn how to configure and manage IIS. You will discover how to configure Web and FTP sites, virtual directories, and IIS security.

After this lesson, you will be able to

- Install IIS
- Set up a Web and FTP site
- Configure a Web default content page
- Create a Web virtual directory
- Modify IIS authentication and security settings

Estimated lesson time: 20 minutes

Installing IIS 6.0

To decrease the attack surface of a Windows Server 2003 system, IIS is not installed by default. It must be added using the Add/Remove Windows Components Wizard from Add Or Remove Programs, located in Control Panel. Select Application Server, click Details, and then select Internet Information Services (IIS). You can control the subcomponents of IIS that are installed, but unless you are very familiar with the role of subcomponents, do not remove any default components. You may, however, want to add components, such as ASP.NET, FTP or FrontPage Server Extensions.

Administering the Web Environment

When IIS is installed, a default Web site is created, allowing you to implement a Web environment quickly and easily. However, you can modify that Web environment to meet your needs. Windows Server 2003 provides the tools necessary to administer IIS and its sites.

After installation has completed, you may open the Internet Information Services (IIS) Manager console from the Administrative Tools group. By default, IIS is configured to serve only static content. To enable dynamic content, select the Web Service Extensions

node. As shown in Figure 6-16, all the extensions are prohibited. Select the appropriate extension and click Allow.

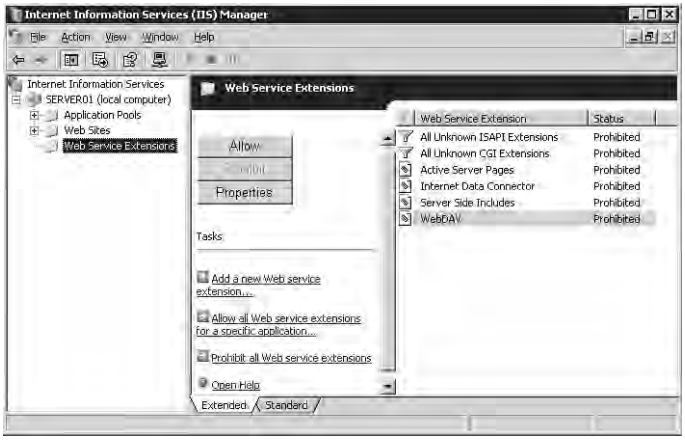

Figure 6-16 The Internet Information Services (IIS) Manager snap-in

The fundamental processes that take place as a client accesses a resource from IIS are

- The client enters a URL (Universal Resource Locator) in either of the following forms:
  - *http://dns.domain.name/virtualdirectory/page.htm* or
  - *ftp://dns.domain.name/virtualdirectory*
- Domain Name Service (DNS) resolves the name to an IP address and returns the address to the client
- The client connects to the server's IP address, using a port that is specific to the service (typically, port 80 for HTTP and port 21 for FTP)
- The URL does not represent the physical path to the resource on the server, but a virtualization of the path. The server translates the incoming request into the phys ical path and produces appropriate resources to the client. For example, the server might list files in the folder to an FTP client, or might deliver the home page to an HTTP client.
- The process can be secured with authentication (credentials, including a user name and password) and authorization (access control through permissions).

You can see this process in action by opening a browser and typing **http://server01**. The server produces the Under Construction page to the client browser.

Configuring and Managing Web and FTP Sites

IIS installation configures a single Web site, the Default Web Site. Although IIS, depending on your server's hardware configuration, can host thousands, or tens of thousands of sites, the Default Web Site is a fine place to explore the functionality and administration of Web sites on IIS. This Web site is accessible if you open a browser and type the URL: **http://server01.contoso.com**. The page that is fetched is the Under Construction page.

Remember that a browser's request to a Web server is directed at the server's IP address, which was resolved from the URL by DNS. The request includes the URL, and the URL often includes only the site name (*www.microsoft.com*, for example). How does the server produce the home page? If you examine the Web Site tab of the Default Web Site Properties, as shown in Figure 6-17, you see that the site is assigned to All Unassigned IP addresses on port 80. So the request from the browser hits port 80 on the server, which then identifies that it is the Default Web Site that should be served.

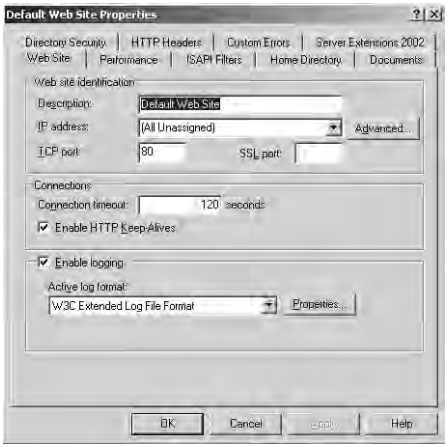

Figure 6-17 The Web Site tab of the Default Web Site Properties dialog box

The next question, then, is what information should be served. If the URL includes only the site name (for example, *www.microsoft.com* or *server01.contoso.com*), then the page that will be returned is fetched from the home directory. The Home Directory tab, as shown in Figure 6-18, displays the physical path to the home directory, typically *c:\inetpub\wwwroot*.

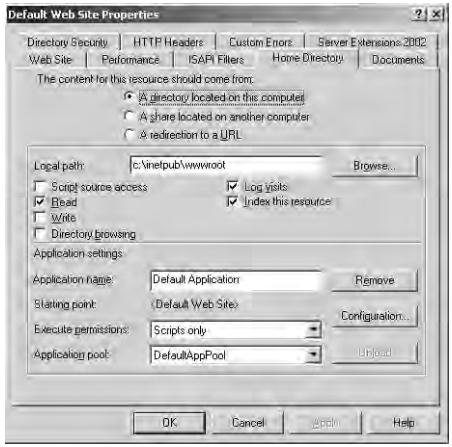

Figure 6-18 The Home Directory tab of the Default Web Site Properties dialog box

Which file, exactly, should be returned to the client? That is defined on the Documents tab, as shown in Figure 6-19. IIS searches for files in the order listed. As soon as it finds a file of that name in the local path of the home directory, that page is returned to the client and the server stops looking for other matches. If no match is found, the IIS returns an error (404–File Not Found) to the client indicating that the page could not be found.

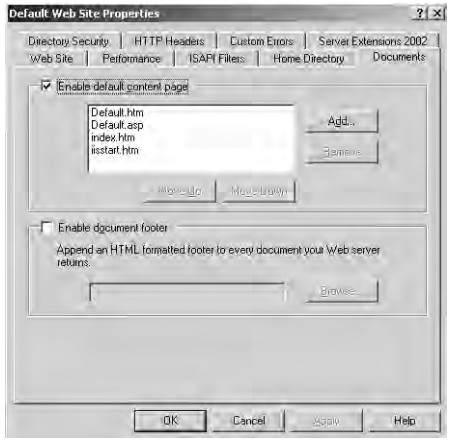

Figure 6-19 The Documents tab of the Default Web Site Properties dialog box

A browser could, of course, refer to a specific page in the URL, for example *http://server01.contoso.com/contactinfo.htm*. In that event, the specific page is fetched from the home directory. If it is not found, a File Not Found error (404) is returned.

To create a Web site, right-click the Web Sites node or an existing Web site in IIS Manager and choose New Web Site. To configure a Web site, open its Properties. You can

configure the IP address of the site. If a server has multiple IP addresses, each IP address can represent a separate Web site. Multiple sites can also be hosted using dif ferent ports for each site, or using host headers. The specifics of these options are beyond the scope of this book. You can also configure the path to the directory that is used as the home directory. And you can modify the list or order of documents that can be fetched as the default content page.

A URL can also include more complex path information, such as *http://www.microsoft.com/windowsserver2003*. This URL is not requesting a specific page; there is no extension such as .htm or .asp on the end of the URL. Instead, it is request ing information from the windowsserver2003 directory. The server evaluates this addi tional component of the URL as a virtual directory. The folder that contains the files referred to as windowsserver2003 can reside anywhere; they do not have to be located on the IIS server.

To create a virtual directory, right-click a Web site and choose New Virtual Directory. The wizard will prompt you for the alias, which becomes the folder name used in the URL, and the physical path to the resource, which can be on a local volume or remote server.

Exam Tip You can also create a Web virtual directory on an NTFS drive by right-clicking a folder, choosing Properties, then clicking the Web Sharing tab.

FTP sites work, and are administered, similarly to Web sites. IIS installs one FTP site, the Default FTP Site, and configures it to respond to all incoming FTP requests (all unassigned addresses, port 21). The FTP site returns to the client a list of files from the folder specified in the Home Directory tab. FTP sites may also include virtual directo ries so that, for example, *ftp://server01.contoso.com/pub* may return resources from a different server than *ftp://server01.contoso.com/vendor-uploads*. FTP URLs and sites do not use default documents.

Complex IIS servers may host tens of thousands of sites, each with customized settings to make them tick. Losing all that configuration information could be painful, so although a normal file system backup might allow you to restore the data files after a failure, the configuration would be lost. To back up or restore IIS configuration, you must back up or restore the *metabase*, an Extensible Markup Language (XML) docu ment that is used to store settings. Right-click the server node in IIS Manager and, from the All Tasks menu, choose Backup/Restore Configuration.

See Also For more information about IIS, see the *Microsoft IIS 6.0 Administrator's Pocket Consultant* (Microsoft Press, 2003).

Securing Files on IIS

Security for files accessed by way of IIS falls into several categories: authentication, authorization through NTFS permissions, and IIS permissions. Authentication is, of course, the process of evaluating credentials in the form of a user name and password. By default, all requests to IIS are serviced by impersonating the user with the IUSR _*computername* account. Before you begin restricting access of resources to specific users, you must create domain or local user accounts and require something more than this default, Anonymous authentication.

Configuring Authentication Methods

You may configure the following authentication methods on the Directory Security tab of the server, a Web (or FTP) site, a virtual directory, or a file:

Web Authentication Options

- **Anonymous authentication** Users may access the public areas of your Web site without a user name or password.
- **Basic authentication** Requires that a user have a local or domain user account. Credentials are transmitted in clear text.
- **Digest authentication** Offers the same functionality as Basic authentication, while providing enhanced security in the way that a user's credentials are sent across the network. Digest authentication relies on the HTTP 1.1 protocol.
- **Advanced Digest authentication** Works only when the user account is part of an Active Directory. Collects user credentials and stores them on the domain controller. Advanced Digest authentication requires the user to be using Internet Explorer 5 or above and the HTTP 1.1 protocol.
- **Integrated Windows authentication** Collects information through a secure form of authentication (sometimes referred to as Windows NT Challenge/Response authentication) where the user name and password are hashed before being sent across the network.
- **Certificate authentication** Adds Secure Sockets Layer (SSL) security through client or server certificates, or both. This option is available only if you have Cer tificate Services installed and configured.
- **.NET Passport authentication** Provides a single sign-in service through SSL, HTTP redirects, cookies, Microsoft JScript, and strong symmetric key encryption.

FTP Authentication Options

- **Anonymous FTP authentication** Gives users access to the public areas of your FTP site without prompting them for a user name or password.
- **Basic FTP authentication** Requires users to log on with a user name and password corresponding to a valid Windows user account.

Defining Resource Access with Permissions

Once authentication has been configured, permissions are assigned to files and folders. A common way to define resource access with IIS is through NTFS permissions. NTFS permissions, because they are attached to a file or folder, act to define access to that resource regardless of how the resource is accessed.

IIS also defines permissions on sites and virtual directories. Although NTFS permissions define a specific level of access to existing Windows user and group accounts, the directory security permissions configured for a site or virtual directory apply to *all* users and groups.

Table 6-2 details Web permission levels:

Table 6-2 IIS Directory Permissions

Permission	Explanation
Read (default)	Users can view file content and properties.
Write	Users can change file content and properties.
Script Source Access	Users can access the source code for files, such as the scripts in an Active Server Pages (ASP) application. This option is available only if either Read or Write permissions are assigned. Users can access source files. If Read permis sion is assigned, source code can be read. If Write permission is assigned, source code can be written to as well. Be aware that allowing users to have read and write access to source code can compromise the security of you server.
Directory browsing	Users can view file lists and collections.

The Execute permissions control the security level of script execution and are as described in Table 6-3.

Table 6-3 Application Execute Permissions

Permission	Explanation
None	Set permissions for an application to None to prevent any programs or scripts from running.
Scripts only	Set permissions for an application to Scripts only to enable applications mapped to a script engine to run in this directory without having permissions set for executables. Setting permissions to Scripts only is more secure than setting them to Scripts and Executables because you can limit the applications that can be run in the directory.
Scripts and Executables	Set permissions for an application to Scripts and Executables to allow any application to run in this directory, including applications mapped to script engines and Windows binaries (.dll and .exe files).

Exam Tip If IIS permissions and NTFS permissions are both in place, the effective permissions will be the more restrictive of the two.

Practice: Administering IIS

In this practice, you will install IIS and configure a new Web site and virtual directory.

Exercise 1: Install IIS

1. Open Add Or Remove Programs from the Control Panel and click Add/Remove Windows Components.
2. Select Application Server and click Details.
3. Select Internet Information Services (IIS) and click Details.
4. Ensure that, at a minimum, Common Files, File Transfer Protocol (FTP) Service, World Wide Web Service, and Internet Information Services Manager are selected.
5. Complete the installation.

Exercise 2: Prepare Simulated Web Content

1. Create a folder on the C:\ drive called ContosoCorp.
2. Open Notepad and create a file with the text "Welcome to Contoso." Save the file as: **"C:\ContosoCorp\Default.htm"** being certain to surround the name with quotation marks.
3. Create a second file with the text "This is the site for Project 101." Save the file as: **"C:\Docs\Project 101\Default.htm"** being certain to surround the name with quotation marks.

Exercise 3: Create a Web Site

1. Open the Internet Information Services (IIS) Manager snap-in from the Administrative Tools group.
2. Right-click the Default Web Site and choose Stop.
3. Right-click the Web Sites node and choose New Web Site.
4. Give the site the name Contoso and the path C:\ContosoCorp. All other default settings are acceptable.

Exercise 4: Create a Secure Virtual Directory

1. Right-click the Contoso site and choose New Virtual Directory.
2. Enter the alias Project101 and the path C:\Docs\Project 101.
3. Open the properties of the Project101 virtual directory.
4. Click Directory Security.
5. In the Authentication and Access Control frame, click Edit.
6. Deselect the option to allow anonymous access. Permission to the files in the site will now require valid user accounts. Click OK twice.
7. Open Internet Explorer and type **http://server01.contoso.com**. The Welcome To Contoso page should appear.
8. Type the URL **http://server01.contoso.com/Project101**. You will be prompted for credentials. Log on as Scott Bishop and the Project101 home page appears.
9. Change the permissions on the C:\Docs\Project 101\Default.htm document so that only Administrators can read the document.
10. Close and reopen Internet Explorer. Connect to *http://server01.contoso.com/Project101* and authenticate as Administrator. The page should appear.
11. Close and reopen Internet Explorer again. Now, connect to the same URL as Scott Bishop. You should receive an Access Denied error (401–Unauthorized).

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're setting up a Web site in IIS on Server01. The site's Internet domain name is *adatum.com*, and the site's home directory is C:\Web\Adatum. Which URL should Internet users use to access files in the home directory of the site?
   a. *http://server01.web.adatum*
   b. *http://web.adatum.com/server01*
   c. *http://server01.adatum/home*
   d. *http://server01.adatum.com*

2. Data for your corporate intranet is currently stored on the D: drive of your IIS server. It is decided that the HR department will serve information about the com pany benefits and policies from its server, and that the URL to access the HR infor mation should be *http://intranet.contoso.com/hr*. What do you need to configure?
   a. A new Web site
   b. A new FTP site
   c. A virtual directory from file
   d. A virtual directory
3. You want to ensure the highest level of security for your corporate intranet without the infrastructure of certificate services. The goal is to provide authentication that is transparent to users, and to allow you to secure intranet resources with the group accounts existing in Active Directory. All users are within the corporate firewall. What authentication method should you choose?
   a. Anonymous Access
   b. Basic Authentication
   c. Digest Authentication
   d. Integrated Windows Authentication

Lesson Summary

- IIS is not installed by default. You can install it using the Windows Components Wizard through Add Or Remove Programs.
- A Web or FTP site's home directory is the physical location of resources to be served by that site.
- A virtual directory is an alias and a path that points the IIS server to the location of resources. The URL takes the form *http://server.dns.name/virtualdirectory*. The resources can be located on a local volume or remote server.
- IIS supports multiple levels of authentication. By default, Anonymous Authentica tion allows any connecting user to access public areas of the site, and Integrated Windows Authentication allows you to assign NTFS permissions to resources that you wish to secure further.
- Access to IIS resources on NTFS volumes is controlled by ACLs, exactly as if the resource were being accessed by the Client For Microsoft Networks.
- IIS has directory and application permissions. If both IIS permissions and NTFS permissions are applied, the more restrictive permissions are effective.

Case Scenario Exercise

Note This Case Scenario exercise is designed to prepare for and to complement the following "Troubleshooting Lab" section. It is recommended that you complete both exercises to gain the maximum learning from these hands-on experiences with Windows Server 2003 file system security.

You must have IIS installed (see Lesson 4, Exercise 1) and have created the group and user accounts as described in this chapter's "Before You Begin" section.

Contoso, Ltd. wants to configure an intranet site for company and departmental news. The specifications call for the site to be easy to use by both employees and the man agers, who will be responsible for updating the news documents. All employees will use the latest version of Internet Explorer to browse the intranet. Managers will use other tools to create Web pages.

Exercise 1: Create Shared Folders and Sample Web Content

Note There are obviously many ways to create and share folders. In this situation, please use the methods described.

1. Open the command prompt.
2. Type the following commands:

```
md c:\ContosoIntranetNews

net share News=c:\ContosoIntranetNews
```

3. Open Notepad and create a file with the text "Contoso Company News." Save the file as **"C:\ ContosoIntranetNews\Default.htm"**, being certain to surround the name with quotation marks.
4. Add the following permission to the C:\ContosoIntranetNews folder:

   Managers: Allow Modify

5. In the C:\ContosoIntranetNews folder's Properties dialog box, click the Web Shar ing tab.
6. From the Share On drop-down list, choose Contoso. If you did not complete the exercises in Lesson 4, you will not have the Contoso Web site; choose the Default Web Site instead. Click Share This Folder and type the alias **News**. The default permissions are adequate. Click OK.

Exercise 2: Optimize Intranet Access

In this exercise, you will confirm the functionality of the intranet and optimize its ease of use.

1. Open Internet Explorer and type the URL: **http://server01.contoso.com/News**.
2. You will be prompted for credentials. Authenticate as Administrator. The Contoso Company News page should appear.
3. Close Internet Explorer.

   You are being prompted for credentials because Company News is not allowing anonymous access. When you create a virtual directory by using the Web Sharing tab, anonymous access is disabled by default.
4. Using IIS manager, open the properties of the News virtual directory.
5. Click the Directory Security tab and click Edit in the Authentication and Access Control frame.
6. Enable anonymous access.
7. Repeat steps 1 through 3 to verify that the change was effective.

Exercise 3: Confirm That Managers Can Modify Intranet Contents

Note To simulate remote management of the intranet contents, it is important that you use the UNC path to the folders and files, as instructed. Do not use a local path.

1. Log off Server01 and log on again as the user Lorrin Smith-Bates, who is a member of the Managers group.
2. Open Notepad and create a document with the text "Good News Contoso!" Save the document as: **"\\server01\news\goodnews.htm"**, being certain to sur round the name in quotation marks and to use the UNC path, not a local path, to the news folder.
3. Are you able to save the file?

   If you followed the instructions of this Case Scenario fully, you should not be able to do so. Continue with the Troubleshooting Lab to identify and solve the problem you just encountered.

Troubleshooting Lab

Note This troubleshooting lab is designed to complement the preceding Case Scenario Exercise. It is recommended that you complete both exercises to gain the maximum learning from these hands-on experiences with Windows Server 2003 file system security.

You must have IIS installed (see Lesson 4, Exercise 1) and have created the group and user accounts as described in this chapter's "Before You Begin" section. You must also have completed at least Exercise 1 of the Case Scenario.

Lorrin Smith-Bates calls the help desk and reports that he is unable to save documents to the intranet news folder. He is creating a Web page in Notepad and saving it to "\\server01\News\goodnews.htm" when the error occurs.

The folder is located at C:\ContosoIntranetNews and is shared as News, and is config ured as a virtual directory, News, for the Contoso Web site. The error message he receives is an Access Denied message. That indicates that his machine is likely able to connect to the server, but that a permission or privilege of some kind prevents him from saving the file.

Log on to Server01 as Administrator to perform these troubleshooting steps.

Step 1: Confirm Group Membership

You are fairly confident that you made Lorrin a member of the Managers group, and that the Managers group has Modify permission to the C:\ContosoIntranetNews folder. How can you confirm Lorrin's group membership?

The Dsget command, discussed in Chapter 3, can enumerate group memberships. Open a command prompt and type the command:

dsget user "CN=Lorrin Smith-Bates,OU=Employees,DC=Contoso,DC=com" -memberof -expand

You should see these groups listed, as well as other groups that may vary depending on which exercises from this book you have completed.

"CN=Managers,OU=Security Groups,DC=contoso,DC=com"

"CN=Project 101 Team,OU=Security Groups,DC=contoso,DC=com"

"CN=Domain Users,CN=Users,DC=contoso,DC=com"

"CN=Print Operators,CN=Builtin,DC=contoso,DC=com"

"CN=Users,CN=Builtin,DC=contoso,DC=com"

How else can you confirm Lorrin's group membership? Open Active Directory Users And Computers and examine the Member Of property page of Lorrin's Properties dia log box.

Step 2: Examine Effective Permissions

Explore the permission assigned to the C:\ContosoIntranetNews folder. You should see, in the Security tab and in the Advanced Security Settings dialog boxes, that Man agers are granted Modify permission.

Click the Effective Permissions tab in the Advanced Security Settings dialog box and select Lorrin's user account. Examine his effective permissions. The permissions should suggest that he is allowed to create files and write data in the folder.

Step 3: Evaluate the Situation

If Lorrin does have effective permissions that allow him to create files and write data, why is he receiving an Access Denied message? If you haven't figured it out already, take a moment to review the Lesson Summaries after Lessons 1 and 4.

The problem might lie in other permissions assigned to the C:\ContosoIntranetNews folder. Share permissions, and Web site or virtual directory permissions define the max imum allowed access, so if one or more of those permissions were configured too restrictively, it could prevent Lorrin from fully using his NTFS Allow Modify permission.

When Lorrin was saving his Web page in Notepad, he was connecting to the server remotely. From the following list, identify the client and the service that were involved:

- FTP Publishing Service
- Worldwide Web Publishing Service
- Telnet Service
- File and Printer Sharing For Microsoft Networks
- Internet browser client
- FTP client
- Telnet client
- Client For Microsoft Networks

Lorrin is using the Client For Microsoft Networks service to connect to Server01's File and Printer Sharing service. You can identify that by examining the path Lorrin speci fied to save the file: "\\server01\News\goodnews.htm." It is a UNC path, which will connect using Microsoft networking.

Knowing that, you can eliminate as a cause of the problem any permissions assigned to the Web site or to the virtual directory; those permissions apply only to connections from Web clients to the Web service.

That leaves one possible cause for permission problems: the Share permissions. The default share permissions in Windows Server 2003 allow the Everyone group only Read permission. Because share permissions define the maximum allowed access, they are overriding the folder's NTFS Allow Modify permission.

Step 4: Solve the Problem

Modify the share permissions on C:\ContosoIntranetNews so that Everyone is allowed Full Control.

Now the business requirements for the intranet news site are that users should only be able to read documents. The default NTFS permission allows users to create files and folders and then, of course, as owners of those files and folders they can do whatever they please.

Lock down NTFS permissions on the folder so that Users have Read & Execute permission, without the special permissions (Create Files/Write Data; Create Folders/Append Data).

Confirm your actions by logging on as Scott Bishop. Scott should be able to see *http://server01.contoso.com/News*. If he connects to \\server01\News, he should *not* be able to create a new file or modify an existing file.

Then log on as Lorrin. Lorrin should also be able to see the intranet news site, but he should also be able to create and modify files in the \\server01\News share. You should be able to create the news document as described in Exercise 3 of the Case Scenario and then access that document at *http://server01.contoso.com/News/goodnews.htm*.

Chapter Summary

- Windows Server 2003 provides new consoles and snap-ins to manage shared folders, audit policy, and IIS. Windows Explorer is still used, as well as the Shared Folder snap-in, to manage NTFS ACLs, although the ACL editor is significantly more powerful.
- NTFS permissions can be allowed or denied; explicit or inherited. A Deny permission takes precedence over an Allow permission; and an explicit permission takes precedence over an inherited permission. The result is that an explicit Allow permission can override an inherited Deny permission.
- Access granted by NTFS permissions may be further restricted by share permissions and IIS permissions on FTP sites, Web sites, virtual directories and documents. Whenever two permission types are assigned to a resource, such as share permissions and NTFS permissions, you must evaluate each set of permissions,

then determine which of the two sets is more restrictive. And that is the set that becomes effective.

- The security descriptor of a file or folder also includes information about the object's owner. The owner, as well as any user with Allow Change permissions, can modify the ACL. Ownership may be assumed by a user with the Allow Take Ownership permission; or may be transferred between users by anyone with the Restore Files And Directories user right.
- The security descriptor also contains auditing entries which, when audit policy is enabled, directs the system to log the specified types of access for the specified users or groups.

Exam Highlights

Before taking the exam, review the key topics and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Familiarize yourself with the tools that are used to configure shared folders, NTFS permissions, auditing and IIS. Spend some time with each snap-in, examining the properties that can be configured, and the role those properties play in managing files and folders.
- Be fluent in the determination of effective permissions: the interaction of explicit, inherited, allowed, and denied permissions for multiple users, groups, computers, and logon types such as Interactive versus Network.
- Know the three steps required to configure auditing, and the strategies you can use to determine what kind of auditing (success or failure) to engage for a particular goal.
- Experience and understand the configuration of a Web site and virtual directory. If you are not experienced with IIS, be certain to implement the Practice in Lesson 4 as well as the Case Scenario and Troubleshooting Lab.

Key Terms

Hidden share A shared folder can be hidden by appending a $ to its share name. Connections can be made to the share using the share's UNC (for example, \\server01\docs$), but the share will not appear on browse lists. Windows Server 2003 creates hidden administrative shares, such as Admin$, Print$, and a hidden share for the root of each disk volume. Only administrators can connect to the hidden administrative shares.

Inheritance By default, permissions assigned to a folder apply to the folder, its subfolders and files. In addition, files and folders are configured by default to allow inheritable permissions from their parent folder or volume to propagate to their ACL. Through these two mechanisms, permissions assigned to a high-level folder are propagated to its contents.

Effective permissions Permissions can be allowed or denied, inherited or explicitly assigned. They can be assigned to one or more users, groups, or computers. The effective permissions are the overall permissions that result and determine the actual access for a security principal.

Ownership Each NTFS file or folder maintains a property that indicates the security principal that owns the resource. The owner is able to modify the ACL of the object at any time, meaning the owner cannot be locked out of the resource. Ownership can be taken and transferred based on the Take Ownership permis sion and the Restore Files And Directories user right, respectively.

The special accounts: Creator Owner, Network, and Interactive These security principals are dynamic, and represent the relationship between a user and a resource. When a user creates a file or folder, they are the Creator Owner of that resource, and any inheritable permissions on the parent folder or volume assigned to Creator Owner will be explicitly assigned to the user on the new object. Network and Interactive represent the connection state of the user—whether the user is connected to the resource from a remote client, or is logged on interactively to the computer that is maintaining the resource.

Audit Object Access policy This policy, available in the Local Security Policy of a standalone Windows Server 2003 computer, or in Group Policy Objects, determines whether access to files, folders, and printers is registered in the Security log. When this policy is enabled, the Auditing Entries for each object determine the types of activities that are logged.

Virtual directory A virtual directory is an IIS object that allows a folder on any local or remote volume to appear as a subfolder of a Web site.

7 Backing Up Data

Exam Objectives in this Chapter:

- Manage backup procedures
  - Verify the successful completion of backup jobs
  - Manage backup storage media
- Configure security for backup operations
- Schedule backup jobs
- Restore backup data

Why This Chapter Matters

You've worked hard to configure and maintain a best practice server environment. You have outfitted the server with a sophisticated RAID subsystem, carefully managed file and share permissions, locked down the server with policy, and physically secured the server to prevent unauthorized interactive log on. But today, none of that matters, because the building's fire sprinklers went off last night, and today your servers are full of water. All that matters today is that you are able to restore your data from backup.

Among the many high priority tasks for any network administrator is the creation and management of a solid backup and restore procedure. Microsoft Windows Server 2003 offers powerful and flexible tools which will enable you to perform backups of local and remote data, including open and locked files, and to schedule those backups for periods of low utilization, such as during the night.

This chapter examines the Ntbackup utility's graphical user interface (GUI) and command-line functionality in the protection of data files. You will learn how to plan an effective backup and media management strategy, how to execute backups, and how to restore data correctly in a variety of scenarios. You will also leverage the new Volume Shadow Copy Service (VSS) to allow faster recovery of data lost by administrators and users alike. Later in the book, we will return to Ntbackup to focus on recovering the operating system during a system restore.

Lessons in this Chapter:

Before You Begin

For hands-on practice using the examples and lab exercises in the chapter, prepare the following:

- Active Directory Users And Computers snap-in
- A Windows Server 2003 (Standard or Enterprise) installed as Server01 and config ured as a domain controller in the domain *contoso.com*

Lesson 1: Fundamentals of Backup

At the core of every backup procedure is a backup tool and a backup plan. Windows Server 2003 provides a robust, flexible utility called Ntbackup. Ntbackup supports much of the functionality found in third-party tools, including the ability to schedule backups, and interacts closely with VSS and the Removable Storage Management (RSM) system. In this lesson, you will examine the conceptual and procedural issues pivotal to the backing up of data, so that you understand the fundamentals of planning for and creating backup jobs with Ntbackup.

After this lesson, you will be able to

- Back up data on local and remote computers
- Understand backup job types
- Create a backup strategy combining normal and incremental or differential backups

Estimated lesson time: 20 minutes

Introducing the Backup Utility

The backup utility in Windows Server 2003, commonly referred to by its executable name, Ntbackup, can be opened by clicking Backup in the Accessories–System Tools program group in the Start menu. Alternatively, it can be launched by typing **ntbackup.exe** in the Run dialog box.

The first time you launch the backup utility, it runs in Wizard mode, as shown in Figure 7-1. This chapter focuses on the more commonly used Backup Utility interface. If you agree with most administrators that it is easier to use the standard utility than the wizard, clear the Always Start In Wizard Mode check box, and then click Advanced Mode.

Figure 7-1 The Backup Or Restore Wizard

As you can see on the utility's Welcome tab in Figure 7-2, you can back up data manually (the Backup tab) or using the Backup Wizard. You can also schedule unattended backup jobs. The Backup Utility is also used to restore data manually (the Restore And Manage Media tab) or using the Restore Wizard. The Automated System Recovery (ASR) Wizard, which backs up critical operating system files, will be discussed later in this book.

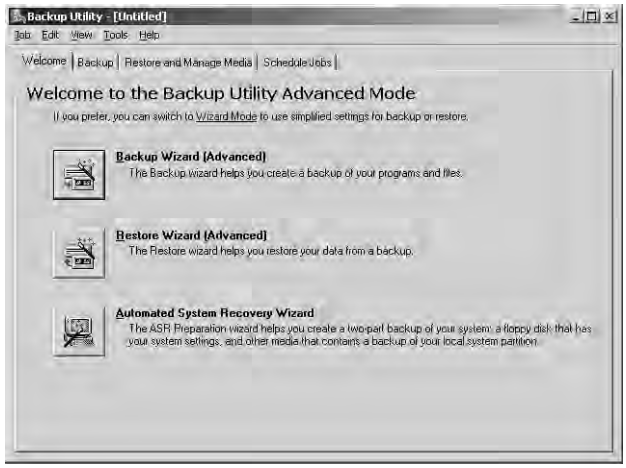

Figure 7-2 The Welcome tab of the Backup Utility

This lesson focuses on data backup planning and execution, and to explore the capability of the Backup Utility we will use the Backup tab, as shown in Figure 7-3, rather than the Backup Wizard.

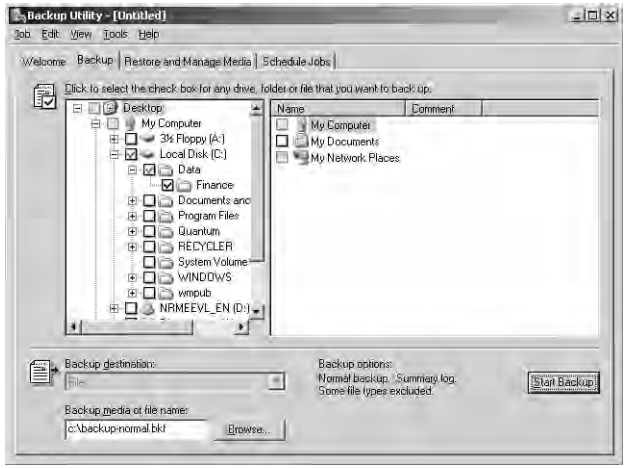

Figure 7-3 The Backup tab of the Backup Utility

Selecting Files to Back Up

You may use the Backup tab to select the files and folders to be backed up. Items may be on local volumes or in network folders. When you select an entire folder for backup, a blue check mark appears. If you select only certain items in a folder, the folder displays a dimmed check mark to indicate a partial backup.

To back up files or folders from remote machines, either select the items from a mapped drive or expand My Network Places. The latter is the equivalent of using a Universal Naming Convention (UNC), such as *Server01**Sharename**Path-to-resource*. Although selecting files and folders through My Network Places is more cumbersome (you must navigate more levels of the interface to locate the files), it has an advantage because drive mappings are more likely to change over time than UNCs.

Tip You can save the set of selected files and folders using the Save Selections command in the Job menu. You can later load the selections using Load Selections from the Job menu, saving the time required to recreate your selection.

Selecting the Backup Destination

Windows Server 2003 allows you to create a backup job on a variety of media types: a tape drive, a removable drive such as the Iomega Jaz drive, and, most importantly, directly to file on a disk volume. If the destination is a tape, the name specified must match the name of a tape that is mounted in the tape device.

If backing up to a file, the Backup Utility creates a .bkf file in the specified location, which can be a local volume or remote folder. It is not uncommon for administrators using the Backup Utility to back up a file on each server and consolidate the resulting files on a central server, which then transfers the backups to removable media. To achieve such a consolidation, the backup destination is configured as either a UNC to a single location on a central server or a local file on each server, which is later copied to a central location.

There are two important limitations of the Backup Utility. First, it does not support writable DVD and CD formats. To work around this limitation, back up to a file, then transfer the file to CD or DVD. Second, backing up to any destination *except* a file requires that the target media be in a device physically attached to the system. This means, for example, that you cannot back up data to a tape drive attached to a remote server.

Determining a Backup Strategy

After selecting the files to back up and specifying the backup destination, there is at least one more critical choice to make. Click Start Backup, then click Advanced, and the Advanced Backup Options dialog box appears, allowing you to specify the backup type. The backup type determines which of your selected files is in fact transferred to the destination media.

Each backup type relates in one way or another to an attribute maintained by every file: archive. The archive (A) attribute is a flag that is set when a file has been created or changed. To reduce the size and duration of backup jobs, most backup types will only transfer to media the files that have their archive attribute set. The most common source of confusion regarding the archive attribute arises from terminology. You will frequently hear, "The file is marked as backed up," which really means that the archive attribute is *cleared* after a particular backup job. The next job will not transfer that file to media. If the file is modified, however, the archive attribute will again be set, and the file will be transferred at the next backup.

Exam Tip As you explore each backup type, keep track of how the archive attribute is used and treated by the backup type. You will need to know the advantages and disadvantages of each backup type and how to fully restore a data structure based on the backup procedures that have been implemented.

Normal Backups

All selected files and folders are backed up. The archive attribute is cleared. A Normal backup does not use the archive attribute to determine which files to back up; all selected items are transferred to the destination media. Every backup strategy begins with a Normal backup that essentially creates a baseline, capturing all files in the backup job.

Normal backups are the most time-consuming and require the most storage capacity of any backup type. However, because they generate a complete backup, normal backups are the most efficient type from which to restore a system. You do not need to restore multiple jobs. Normal backups clear the archive attribute from all selected files.

Incremental Backups

Selected files with the archive attribute set are backed up. The archive attribute is cleared. Selected files with the archive flag are transferred to the destination media, and the flag is cleared. If you perform an incremental backup one day after a normal backup has been performed, the job will contain only the files that were created or changed during that day. Similarly, if you perform an incremental backup one day after another incremental backup, the job will contain only the files that were created or changed during that day.

Incremental backups are the fastest and smallest type of backup. However they are less efficient as a restore set, because you must restore the normal backup and then restore, in order of creation, each subsequent incremental backup.

Differential Backups

Selected files with the archive attribute set are backed up. The archive attribute is not cleared. Because a differential backup uses the archive attribute, the job includes only files that have been created or changed since the last normal or incremental backup. A differential backup does not clear the archive attribute; therefore, if you perform differential backups two days in a row, the second job will include all the files in the first backup, as well as any files that were created or changed during the second day. As a result, differential backups tend to be larger and more time-consuming than incremental backups, but less so than normal backups.

Differential backups are significantly more efficient than incremental backups as a restore set, however. To fully restore a system you would restore the normal backup and the most recent differential backup.

Copy Backups

All selected files and folders are backed up. Copy neither uses nor clears the archive attribute. Copy backups are not used for typical or scheduled backups. Instead, copy backups are useful to move data between systems or to create an archival copy of data at a point in time without disrupting standard backup procedures.

Daily Backups

All selected files and folders that have changed during the day are backed up, based on the files' modify date. The archive attribute is neither used nor cleared. If you want to back up all files and folders that change during the day without affecting a backup schedule, use a daily backup.

Combining Backup Types

Although creating a normal backup every night ensures that a server can be restored from a single job the next day, a normal backup may take too much time to create, perhaps causing the overnight job to last well into the morning, thus disrupting performance during working hours. To create an optimal backup strategy, you must take into account the time and size of the backup job, as well as the time required to restore a system in the event of failure. Two common solutions are:

- **Normal and differential backups** On Sunday a normal backup is performed, and on Monday through Friday nights, differential backups are performed. Differential backups do not clear the archive attribute, which means that each backup includes all changes since Sunday. If data becomes corrupt on Friday, you only

need to restore the normal backup from Sunday and the differential backup from Thursday. This strategy takes more time to back up, particularly if data changes frequently, but is easier and faster to restore, because the backup set is on fewer disks or tapes.

- **Normal and incremental backups** On Sunday a normal backup is performed, and on Monday through Friday incremental backups are performed. Incremental backups clear the archive attribute, which means that each backup includes only the files that changed since the previous backup. If data becomes corrupt on Friday, you need to restore the normal backup from Sunday and each of the incremental backups, from Monday through Friday. This strategy takes less time to back up but more time to restore.

Practice: Performing Different Backup Types

In this practice, you will create several backup jobs, examining the role of the archive attribute.

Exercise 1: Create Sample Data

1. Open Notepad and create a text file with the following lines. Type each line carefully.

```
md c:\Data
net share data=C:\Data
md c:\Data\Finance
cd c:\data\Finance
echo Historical Financial Data > Historical.txt
echo Current Financials > Current.txt
echo Budget > Budget.txt
echo Financial Projections > Projections.txt
```

2. Save the file as "c:\createfiles.bat" including the quotation marks.
3. Open the command prompt and type **cd c:**.
4. Type the command **createfiles.bat**.
5. Open Windows Explorer and navigate to the c:\data\finance directory. You should see the following display:

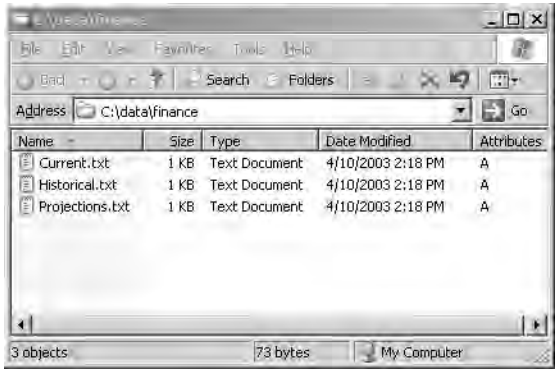

6. If the Attributes column is not visible, right-click the column headers Date Modified and select Attributes. The archive attribute is displayed.

Note Leave Windows Explorer open on C:\Data\Finance. You will refer to it throughout this practice.

Exercise 2: Perform a Normal Backup

1. Open the Backup Utility by running Ntbackup.exe from the command line, or selecting Backup from the Accessories–System Tools group on the Start menu.
2. Clear the Always Start In Wizard Mode check box.
3. Click Advanced Mode.
4. Select the Backup tab.
5. Expand My Computer, the C drive, and then the Data folder so that you can select the Finance folder.

   The Finance folder has a blue check mark, meaning a complete backup, whereas its parent folder has a dimmed check mark, indicating a partial backup. Any files *added* to the Finance folder will be included in the backup, but any files added to the Data folder will not.
6. On the Job menu, choose Save Selections.
7. Save the selections as **Finance Backup.bks**.
8. In the Backup Media Or Filename box, type **c:\backup-normal.bkf**.

Note In production environments you will be likely to use removable media for backups, but to keep hardware requirements to a minimum, practices in this lesson will back up and restore using local files. If you have access to a tape drive, feel free to use it during these practices.

9. Click Start Backup and then click Advanced.
10. Confirm that Normal is selected in the Backup Type drop-down box, and then click OK.
11. Select Replace The Data On The Media With This Backup and click Start Backup.
12. Observe the Backup Progress dialog box. When the backup is complete, click Report.
13. Examine the report. No errors should be reported.
14. Close the report and the Backup Utility.

    Note that in Windows Explorer, the Attributes column no longer shows the archive attribute.

Exercise 3: Perform Differential Backups

1. Open C:\Data\Finance\Current.txt and add some text. Save and close the file.
2. Examine C:\Data\Finance in Windows Explorer. What files are showing the archive attribute?

   Only the one you just changed.
3. Open the Backup Utility and click the Backup tab.
4. From the Job menu, choose Load Selections to load Finance Backup selections.
5. In the Backup Media Or Filename box, type **c:\backup-diff-day1.bkf**.
6. Click Start Backup.
7. Click Advanced and select Differential as the backup type.
8. Start the backup and, when complete, confirm that no errors occurred.
9. Close the Backup Utility.
10. Examine the folder in Windows Explorer. Which files have their archive attribute set?

    The file Current.txt is still flagged for archiving.
11. Open the Budget file and make some changes. Save and close the file. Confirm that its archive attribute is now set.
12. Repeat steps 3 through 9, creating a backup job in the location: **c:\backup-diff-day2.bkf**. Be sure to look at the resulting backup report. How many files were copied for the backup?

    Two.

Exercise 4: Perform Incremental Backups

1. Open the Backup Utility and click the Backup tab.
2. From the Job menu, choose Load Selections to load Finance Backup selections.
3. In the Backup Media Or Filename box, type **c:\backup-inc-day2.bkf**.
4. Click Start Backup.
5. Click Advanced and select Incremental as the backup type.
6. Start the backup and, when complete, confirm that no errors occurred.
7. Close the Backup Utility.
8. Examine the folder in Windows Explorer. Which files have their archive attribute set?

   None.

9. Open the Projections file and make some changes. Save and close the file. It should show the archive attribute in Windows Explorer.
10. Repeat steps 1 through 8, creating a backup job in the location: **c:\backup-inc-day3.bkf**.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Which of the following locations are *not* allowed to be used for a backup of a Windows Server 2003 system?
   a. Local tape drive
   b. Local CD-RW
   c. Local hard drive
   d. Shared folder on a remote server
   e. Local DVD+R
   f. Local removable drive
   g. Tape drive on a remote server
2. You are to back up a Windows Server 2003 file server every evening. You perform a manual, normal backup. You will then schedule a backup job to run every evening for the next two weeks. Which backup type will complete the fastest?
   a. Normal
   b. Differential
   c. Incremental
   d. Copy
3. You are to back up a Windows Server 2003 file server every evening. You perform a manual, normal backup. You will then schedule a backup job to run every evening for the next two weeks. Which backup type will provide the simplest recovery of lost data?
   a. Normal
   b. Differential
   c. Incremental
   d. Daily

4. You are to back up a Windows Server 2003 file server every evening. You perform a normal backup. On the second evening, you consider whether to use incremental or differential backup. Will there be any difference in the speed or size of those two backup jobs? If the server were to fail the following day, would there be any difference in the efficiency of recovery?

5. Review the steps taken during the Practice. Predict the contents of the following backup jobs:
   - ❑ backup-normal.bkf
   - ❑ backup-diff-day1.bkf
   - ❑ backup-diff-day2.bkf
   - ❑ backup-inc-day2.bkf
   - ❑ backup-inc-day3.bkf

   Are there any differences between the contents of backup-diff-day2 and backup-inc-day2?

Note You can find the answers in the Questions and Answers section at the end of the lesson. However, you should test your predictions by performing the Practice in Lesson 2.

Lesson Summary

- The Backup Utility, Ntbackup, allows you to back up and restore data from local and remote folders.
- You may back up to local files, tape drives, and removable media or to shared folders on remote servers. You cannot back up to writable CD or DVD formats.
- A normal backup is a complete backup of all selected files and folders. It is always the starting point of any backup strategy.
- An incremental backup copies selected files that have changed since the most recent normal or incremental backup. Both normal and incremental backups clear the archive attribute.

- A differential backup copies all selected files that have changed since the last normal or incremental backup. Differential backups do *not* clear the archive attribute.
- Copy backups and daily backups are less frequently used. They back up all selected files, in the case of Copy backup, or files modified on a specific date, in the case of Daily backup. They do not reset the archive attribute, so they can be used to capture data for backup or transfer without interfering with the normal backup schedule.

Lesson 2: Restoring Data

In conjunction with the design of a backup strategy, you must create and verify restore procedures to ensure that appropriate personnel are knowledgeable in the concepts and skills that are critical to data recovery. This lesson will share the processes and options available for restoring data using the Backup Utility.

After this lesson, you will be able to

- Restore data to its original location or an alternate folder
- Configure restore options

Estimated lesson time: 10 minutes

Restoring with the Backup Utility

Restoring data is a straightforward procedure. After opening the Backup Utility and clicking the Restore And Manage Media tab as shown in Figure 7-4, you will be able to select the backup set from which to restore. Windows Server 2003 will then display the files and folders that the backup set contains by examining the backup set's catalog. You can then select the specific files or folders you wish to restore. As with the backup selection, a blue check mark indicates that a file or folder will be fully restored. A dimmed check mark on a folder means that some, but not all, of its contents will be restored.

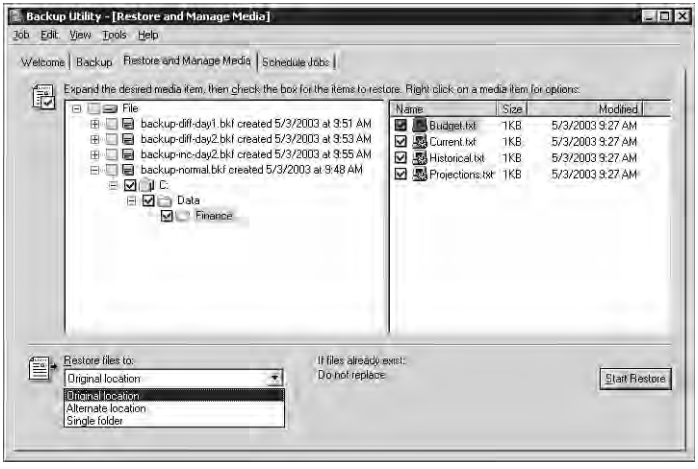

Figure 7-4 The Backup Utility's Restore And Manage Media tab

You are also asked to specify the restore location. For this option, you have three choices:

- **Original location** Files and folders will be restored to the location from which they were backed up. The original folder structure will be maintained or, if folders were deleted, re-created.
- **Alternate location** Files and folders will be restored to a folder you designate in the Alternate Location box. The original folder structure is preserved and cre ated beneath that folder, where the designated alternate location is equivalent to the root (volume) of the backed up data. So, for example, if you backed up a folder C:\Data\Finance and you restored the folder to C:\Restore, you would find the Finance folder in C:\Restore\Data\Finance.
- **Single folder** Files are restored to the folder you designate, but the folder struc ture is not maintained. All files are restored to a single folder.

After selecting the files to restore and the restore location, click Start Restore. Click OK and the restore process will begin. Confirm that no errors occurred.

Restore Options

Windows Server 2003 supports several options for how files in the restore location are handled during a restore. The following options are found in the Backup Utility's Tools–Options command, on the Restore tab shown in Figure 7-5:

- **Do Not Replace The File On My Computer.** This option, the default, causes the Restore utility to skip files that are already in the target location. A common scenario leading to this choice is one in which some, but not all, files have been deleted from the restore location. This option will restore such missing files with the backed-up files.
- **Replace The File On Disk Only If The File On Disk Is Older.** This option directs the restore process to overwrite existing files unless those files are more recent than the files in the backup set. The theory is that if a file in the target loca tion is more recent than the backed-up copy, it is possible that the newer file con tains information that you do not want to overwrite.
- **Always Replace The File On My Computer.** Under this restore option, all files are overwritten by their backed-up versions, regardless of whether the file is more recent than the backup. You will lose data in files that were modified since the backup date. Any files in the target location that are *not* in the backup set will remain, however.

After selecting files to restore, restore options and a restore destination, click Start Restore, and then confirm the restore. The Start Restore dialog box appears.

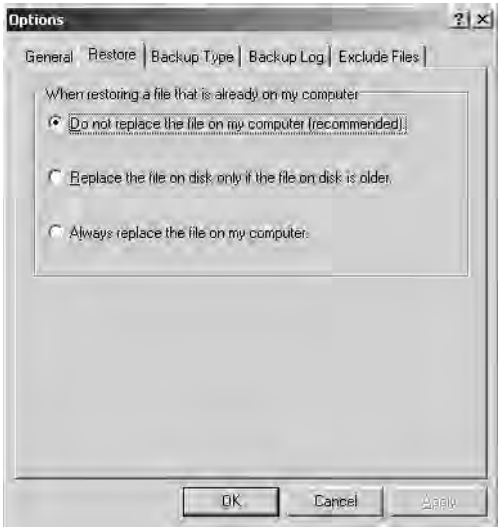

Figure 7-5 Restore tab options

Before confirming the restore, you can configure how the restore operation will treat security settings on the backed-up files by clicking Advanced in the Confirm Restore dialog box and selecting the Restore Security option. If data was backed up from, and is being restored to, an NTFS volume, the default setting will restore permissions, audit settings, and ownership information. Deselecting this option will restore the data without its security descriptors, and all restored files will inherit the permissions of the target restore volume or folder.

Practice: Restoring Data

In this practice, you will verify your backup and restore procedures using a common method: restoring to a test location.

Exercise 1: Verify Backup and Restore Procedures

To verify backup and restore procedures, many administrators will perform a test restore of a backup set. So as not to damage production data, that test restore is targeted not at the original location of the data, but at another folder, which can then be discarded following the test. In a production environment, your verification should include restoring the backup to a "standby" server, which would entail making sure that the backup device (that is, the tape drive) is correctly installed on a server that can host data in the event that the primary server fails. To do this, perform the following steps:

1. Open the Backup Utility.
2. Click Restore And Manage Media.
3. Click the plus sign to expand the file.

4. Click the plus sign to expand Backup-normal.bkf.
5. Click the check box to select C:.
6. Expand C:, Data, and Finance. You will notice that your selection of the C: folder has selected its child folders and files.
7. In the Restore Files To drop-down box, select Alternate Location.
8. In the Alternate Location field, type **C:\TestRestore**.
9. Click Start Restore.
10. In the Confirm Restore dialog box, click OK.
11. When the restore job is complete, click Report and examine the log of the restore operation.
12. Open the C:\TestRestore folder and verify that the folder structure and files restored correctly.
13. Repeat steps 1 through 10, this time restoring the file backup-diff-day2.bkf. When the restore job is finished, continue to step 14 to examine its report.
14. When the restore job finishes, click Report to view the restore job log. If you acci dentally close the job status window, choose the Report command from the Tools menu, select the most recent report and click View.
15. Examine the report for the job you just restored. How many files were restored?

    None.

    Why?

    The answer lies in the restore options.

16. Choose the Options command from the Tools menu and click the Restore tab. Now you can identify the problem. The default configuration of the backup utility is that it does not replace files on the computer. Therefore, the differential job, which contains files that were updated after the normal backup, was not successfully restored.
17. Choose Always Replace The File On My Computer.
18. Repeat the restore operation of backup-diff-day2.bkf. The report should confirm that two files were restored.
19. You have now verified your backup and restore procedures, including the need to modify restore options. Delete the C:\TestRestore folder.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. A user has accidentally deleted the data in a Microsoft Word document and saved the document, thereby permanently altering the original file. A normal backup operation was performed on the server the previous evening. Which restore option should you select?
   - **a.** Do Not Replace The File On My Computer.
   - **b.** Replace The File On Disk Only If The File On Disk Is Older.
   - **c.** Always Replace The File On My Computer.
2. An executive has returned from a business trip. Before the trip, she copied files from a network folder to her hard drive. The folder is shared with other executives, who modified their files in the folder while she was away. When she returned, she moved her copy of the files to the network share, thereby updating her files with the changes she made while away, but also overwriting all the files that had been changed by other executives. The other executives are unhappy that their files have been replaced with the versions that were active when she left for her trip. Luckily, you performed a normal backup operation on the folder the previous evening. What restore option should you choose?
   - **a.** Do Not Replace The File On My Computer.
   - **b.** Replace The File On Disk Only If The File On Disk Is Older.
   - **c.** Always Replace The File On My Computer.
3. You would like to test the restore procedures on your server, but would also like to avoid affecting the production copies of the backed-up data. What is the best restore location to use?
   - **a.** Original location
   - **b.** Alternate location
   - **c.** Single folder

Lesson Summary

- The Backup Utility will also allow you to restore backed-up data.
- When restoring a lost file or folder, it is common to select Original Location as the restore location.
- When testing restore procedures, it is common to select Alternate Location as the restore location so that you do not affect the original copies of the backed-up files and folders.
- When restoring a differential or incremental backup set after restoring the normal backup set, you will need to select the restore option Always Replace The File On My Computer.
- When restoring a folder in which files have been lost, but some files are intact, you should select the restore option Do Not Replace The File On My Computer or Replace The File On Disk Only If The File On Disk Is Older.

Lesson 3: Advanced Backup and Restore

Now that you have created a backup plan and verified your procedures for backup and restore, you will want to understand the process in more depth so that you can config ure backup operations to be more flexible, more automated or perhaps even easier. This lesson will explore the technologies underlying data backup, such as VSC and RSM, and will lay out options for scripting and scheduling backup operations. You will then leverage the new Shadow Copies Of Shared Folders feature to enable users to recover from simple data loss scenarios without administrative intervention.

After this lesson, you will be able to

- Configure group membership to enable a user to perform backup and restore operations
- Manage tape backup media
- Catalog backup sets
- Configure backup options
- Execute a backup from the command prompt
- Schedule backup jobs
- Configure and utilize Shadow Copies Of Shared Folders

Estimated lesson time: 30 minutes

Understanding VSS

Windows Server 2003 offers VSS, also referred to as "snap backup." VSS allows the backing up of databases and other files that are held open or locked due to operator or system activity. Shadow copy backups allow applications to continue to write data to a volume during backup, and allow administrators to perform backups at any time without locking out users or risking skipped files.

Although VSS is an important enhancement to the backup functionality of Windows Server 2003, it is nevertheless best practice to perform backups when utilization is low. If you have applications that manage storage consistency differently while files are open, that can affect the consistency of the files in the backup of those open files. For critical applications, or for applications such as Microsoft SQL Server that offer native backup capabilities, consult the documentation for the application to determine the recommended backup procedure.

Backup Security

You must have the Backup Files And Directories user right, or NTFS Read permission, to back up a file. Similarly, you must have the Restore Files And Directories user right, or NTFS Write permission to the target destination, to restore a file. Privileges are

assigned to both the Administrators and Backup Operators groups, so the minimum required privileges can be given to a user, a group, or a service account by nesting the account in the Backup Operators group on the server.

Users with the Restore Files And Directories user right can remove NTFS permissions from files during restore. In Windows Server 2003, they can additionally transfer own ership of files between users.

Therefore, it is important to control the membership of the Backup Operators group and to physically secure backup tapes. A "loose" backup tape makes it easy for any intelligent individual to restore and access sensitive data.

Managing Media

The Backup Utility of Windows Server 2003 works closely with the RSM service. RSM, which is designed to manage robotic tape libraries and CD-ROM libraries, accepts requests for media from other services or, in this case, applications, and ensures that the media is correctly mounted or loaded.

RSM is also used with single-media devices, such as a manually loaded backup tape drive, CD-ROM, or Iomega Jaz drive. In the case of single-media drives, RSM keeps track of media through their labels or serial numbers. The impact of RSM is that, even in a single-media drive backup system, each tape must have a unique label.

Media Pools

The Backup Utility of Windows Server 2003 manages tapes with RSM using *media pools*, as seen in Figure 7-6.

Figure 7-6 Media pools

There are four media pools related to backup:

- **Unrecognized** Tape media that are completely blank or in a foreign format are contained in the Unrecognized pool until they are formatted.

- **Free** This pool contains newly formatted tape media, as well as tapes that have been specifically marked as free by an administrator. Free media can be moved into the backup media pool by writing a backup set to them.
- **Backup** This pool contains media that have been written to by the Backup Utility. The Backup Utility will only write to media in the Free media pool (and it will label the tape with the name you enter just before starting the backup) and to media, specified by name, in the Backup media pool.
- **Import** This pool contains tape media that are not cataloged on the local disk drive. Cataloging such a tape will move the tape into the backup media pool.

Managing Tapes and Media Pools

In conjunction with backup procedures and tape rotation, you will need to manage your tapes in and out of these media pools. To that end, the following actions are available from the Restore And Manage Media page of the Backup Utility:

- **Format a tape** Right-click a tape and choose Format. Formatting is not a secure way to erase tapes. If you need to erase tapes for legal or security reasons, use an appropriate third-party utility. Formatting does, however, prepare a tape and move it into the free media pool. Not all drives support formatting.
- **Retension a tape** Right-click a tape and choose Retension. Not all drives support retensioning.
- **Mark a tape as free** Right-click a tape and choose Mark As Free. This moves the tape into the free media pool. It does *not* erase the tape. If you need to erase tapes for legal reasons, use an appropriate third-party utility.

Catalogs

When the Backup Utility creates a backup set, it also creates a catalog listing files and folders included in the backup set. That catalog is stored on the disk of the server (the local or on-disk catalog) and in the backup set itself (the on-media catalog). The local catalog facilitates quick location of files and folders to restore. The Backup Utility can display the catalog immediately, rather than load the catalog from the typically slower backup media. The on-media catalog is critical if the drive containing the local catalog has failed, or if you transfer the files to another system. In those cases, Windows can recreate the local catalog from the on-media catalog.

The Restore And Manage Media page of the Backup Utility allows you to manage catalogs, as follows:

- **Delete Catalog** Right-click a backup set and choose Delete Catalog if you have lost or damaged the backup media or if you are transferring files to another system and no longer require its local catalog. The on-media catalog is not affected by this command.

- **Catalog** A tape from a foreign system that is not cataloged on the local machine will appear in the import media pool. Right-click the media and choose the Cata log command. Windows will generate a local catalog from the tape or file. This does not create or modify the on-media catalog.

Tip If you have all the tapes in the backup set and the tapes are not damaged or corrupted, open the backup Options dialog box and, on the General tab, select Use The Catalogs On The Media To Speed Up Building Restore Catalogs On Disk. If you are missing a tape in the backup set or a tape is damaged or corrupted, clear that option. This will ensure that the catalog is complete and accurate; however, it might take a long time to create the catalog.

Backup Options

Backup options are configured by choosing the Options command from the Tools menu. Many of these options configure defaults that are used by the Backup Utility and the command-line backup tool, Ntbackup. Those settings can be overridden by options of a specific job.

General Options

The General tab of the Options dialog box includes the following settings:

- **Compute Selection Information Before Backup And Restore Operations** Backup estimates the number of files and bytes that will be backed up or restored before beginning the operation.
- **Use The Catalogs On The Media To Speed Up Building Restore Catalogs On Disk** If a system does not have an on-disk catalog for a tape, this option allows the system to create an on-disk catalog from the on-media catalog. However, if the tape with the on-media catalog is missing or if media in the set is damaged, you can deselect this option and the system will scan the entire backup set (or as much of it as you have) to build the on-disk catalog. Such an operation can take several hours if the backup set is large.
- **Verify Data After The Backup Completes** The system compares the contents of the backup media to the original files and logs any discrepancies. This option obviously adds a significant amount of time for completing the backup job. Discrep ancies are likely if data changes frequently during backup or verification, and it is not recommended to verify system backups because of the number of changes that happen to system files on a continual basis. So long as you rotate tapes and discard tapes before they are worn, it should not be necessary to verify data.

- **Backup The Contents Of Mounted Drives** A mounted drive is a drive volume that is mapped to a folder on another volume's namespace, rather than, or in addition to, having a drive letter. If this option is deselected, only the path of the folder that is mounted to a volume is backed up; the contents are not. By selecting this option, the contents of the mounted volume is also backed up. There is no disadvantage in backing up a mount point, however if you back up the mount point and the mounted drive as well, your backup set will have duplication.

If you primarily back up to file and then save that file to another media, *clear* the following options. If you primarily back up to a tape or another media managed by Removable Storage, *select* the following options.

- Show Alert Message When I Start the Backup Utility And Removable Storage Is Not Running.
- Show Alert Message When I Start The Backup Utility And There Is Recognizable Media Available.
- Show Alert Message When New Media Is Inserted.
- Always Allow Use Of Recognizable Media Without Prompting.

Tip The Always Allow Use Of Recognizable Media Without Prompting option can be selected if you are using local tape drives for backup only, not for Remote Storage or other functions. The option eliminates the need to allocate free media using the Removable Storage node in the Computer Management console.

Backup Logging

The Options dialog has a tab called Backup Log. Logging alerts you to problems that might threaten the viability of your backup, so consider your logging strategy as well as your overall backup plan. Although detailed logging will list every file and path that was backed up, the log is so verbose you are likely to overlook problems. Therefore, summary logging is recommended, and is the default. Summary logs report skipped files and errors.

The system will save 10 backup logs to the path *%UserProfile%*\Local Settings\Application Data\Microsoft\Windows NT\Ntbackup\Data. There is no way to change the path or the number of logs that are saved before the oldest log is replaced. You can, of course, include that path in your backup and thereby back up old logs.

File Exclusions

The Exclude Files tab of the Options dialog box also allows you to specify extensions and individual files that should be skipped during backup. Default settings result in the Backup Utility's skipping the page file, temporary files, client-side cache, debug folder, and the File Replication Service (FRS) database and folders, as well as other local logs and databases.

Files can be excluded based on ownership of the files. Click Add New under Files Excluded For All Users to exclude files owned by any user. Click Add New under Files Excluded For User <*username*> if you want to exclude only files that you own. You can specify files based on Registered File Type or based on an extension using the Cus tom File Mask. Finally, you can restrict excluded files to a specific folder or hard drive using the Applies To Path and the Applies To All Subfolders options.

Advanced Backup Options

After selecting files to back up, and clicking Start Backup, you can configure additional, job-specific options by clicking Advanced. Among the more important settings are the following:

- **Verify Data After Backup** This setting overrides the default setting in the Backup Options dialog box.
- **If Possible, Compress The Backup Data To Save Space** This setting com presses data to save space on the backup media, an option not available unless the tape drive supports compression.
- **Disable Volume Shadow Copy** VSS allows the backup of locked and open files. If this option is selected, some files that are open or in use may be skipped.

The Ntbackup Command

The Ntbackup command provides the opportunity to script backup jobs on Windows Server 2003. Its syntax is

```
Ntbackup backup {"path to backup" or "@selectionfile.bks"} /j "Job Name" options
```

The command's first switch is *backup*, which sets its mode—you cannot restore from the command line. That switch is followed by a parameter that specifies what to back up. You can specify the actual path to the local folder, network share, or file that you want to back up. Alternatively, you can indicate the path to a backup selection file (.bks file) to be used with the syntax @*selectionfile*.bks. The at (@) symbol must precede the name of the backup selection file. A backup selection file contains information on

the files and folders you have selected for backup. You have to create the file using the graphical user interface (GUI) version of the Backup Utility.

The third switch, /J "*JobName*", specifies the descriptive job name, which is used in the backup report.

You can then select from a staggering list of switches, which are grouped below based on the type of backup job you want to perform.

Backing Up to a File

Use the switch

/F "*FileName*"

where *FileName* is the logical disk path and file name. You must not use the following switches with this switch: /T /P /G.

The following example backs up the remote Data share on Server01 to a local file on the E drive:

```
ntbackup backup "\\server01\Data" /J "Backup of Server 01 Data folder" /F
"E:\Backup.bkf"
```

Appending to a File or Tape

Use the switch:

/A

to perform an append operation. If appending to a tape rather than a file, you must use either /G or /T in conjunction with this switch. Cannot be used with /N or /P.

The following example backs up the remote Profiles share on Server02 and appends the set to the job created in the first example:

```
ntbackup backup "\\server02\Profiles" /J "Backup of Server 02 Profiles folder" /F
"E:\Backup.bkf" /A
```

Backing Up to a New Tape or File, or Overwriting an Existing Tape

Use the switch:

/N "*MediaName*"

where *MediaName* specifies the new tape name. You must not use /A with this switch.

Backing Up to a New Tape

Use the switch

/P "*PoolName*"

where *PoolName* specifies the media pool that contains the backup media. This is usu ally a subpool of the backup media pool, such as 4mm DDS. You cannot use the /A, /G, /F, or /T options if you are using /P.

The following example backs up files and folders listed in the backup selection file c:\backup.bks to a tape drive:

```
ntbackup backup @c:\backup.bks /j "Backup Job 101" /n "Command Line Backup Job" /p
"4mm DDS"
```

Backing Up to an Existing Tape

To specify a tape for an append or overwrite operation, you must use either the /T or /G switch along with either /A (append) or /N (overwrite). Do not use the /P switch with either /T or /G.

To specify a tape by name, use the /T switch with the following syntax:

/T "*TapeName*"

where *TapeName* specifies a valid tape in the media pool.

To back up the selection file and append it to the tape created in the previous example, you would use this command line:

```
ntbackup backup @c:\backup.bks /j "Backup Job 102" /a /t "Command Line Backup Job"
```

To specify a tape by its GUID, rather than its name, use the /G switch with the follow ing syntax:

/G "*GUIDName*"

where *GUIDName* specifies a valid tape in the media pool.

Job Options

For each of the job types described above, you can specify additional job options using these switches:

- **/M {*BackupType*}** Specifies the backup type, which must be one of the follow ing: normal, copy, differential, incremental, or daily.
- **/D {"*SetDescription*"}** Specifies a label for the backup set.

- **/V:{yes | no}** Verifies the data after the backup is complete.
- **/R:{yes | no}** Restricts access to this tape to the owner or members of the Administrators group.
- **/L:{f | s | n}** Specifies the type of log file: f=full, s=summary, n=none (no log file is created).
- **/RS:{yes | no}** Backs up the migrated data files located in Remote Storage.

Tip The /RS command-line option is not required to back up the local Removable Storage database, which contains the Remote Storage placeholder files. When you backup the %*Systemroot*% folder, Backup automatically backs up the Removable Storage database as well.

- **/HC:{on | off}** Uses hardware compression, if available, on the tape drive.
- **/SNAP:{on | off}** Specifies whether the backup should use a Volume Shadow Copy.

Scheduling Backup Jobs

To schedule a backup job, create the job in the Backup Utility then click Start Backup and configure advanced backup options. After all options have been configured, click Schedule and, in the Set Account Information dialog box, type the user name and password of the account to be used by the backup job.

Security Alert Security best practices suggest that you create an account for each service, rather than run services under the System account. Do not configure a service to run using a User account, such as your User account or the Administrator account. When the password changes on a User account, you must modify the password setting on all services that run under the context of that account. The account for the backup job should belong to the Backup Operators group.

In the Scheduled Job Options dialog box, enter a job name and click Properties. The Schedule Job dialog box appears, as shown in Figure 7-7. Configure the job date, time, and frequency. The Advanced button will let you configure additional schedule settings including a date range for the job. The Settings tab of the Schedule Job dialog box allows you to refine the job, for example, by specifying that the job should only take place if the machine has been idle for a period of time.

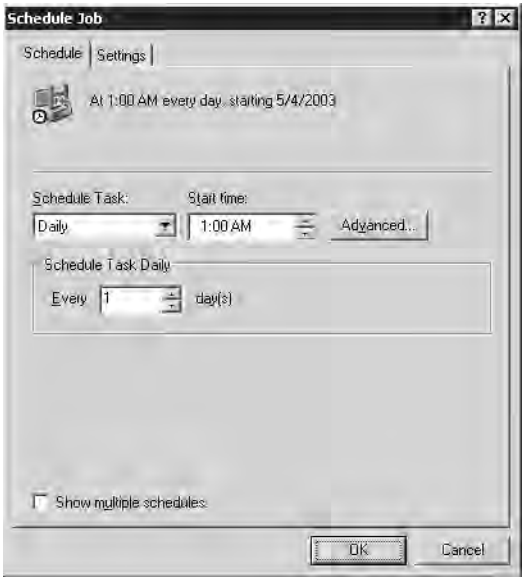

Figure 7-7 The Schedule Job dialog box

Once a job has been scheduled, you can edit the schedule by clicking the Schedule Jobs tab of the Backup Utility. Jobs are listed on a calendar. Click a job to open its schedule. Although you can also add a backup job by clicking Add Job on the Sched ule Jobs tab, clicking Add Job will launch the backup wizard so that you can select the files to back up and some of the properties of the backup job. Most administrators find it more convenient to create a backup job on the Backup tab directly, then click Start Backup and Schedule, as described above.

Shadow Copies of Shared Folders

Windows Server 2003 supports another way for administrators and users alike to recover quickly from damage to files and folders. Using VSS, Windows Server 2003 automatically caches copies of files as they are modified. If a user deletes, overwrites, or makes unwanted changes to a file, you can simply restore a previous version of the file. This is a valuable feature, but is not intended to replace backups. Instead, it is designed to facilitate quick recovery from simple, day-to-day problems—not recovery from significant data loss.

Enabling and Configuring Shadow Copies

The Shadow Copies feature for shared folders is not enabled by default. To enable the feature, open the Properties dialog box of a drive volume from Windows Explorer or the Disk Management snap-in. On the Shadow Copies tab, as shown in Figure 7-8, select the volume and click Enable. Once enabled, all shared folders on the volume

will be shadowed; specific shares on a volume cannot be selected. You can, however, manually initiate a shadow copy by clicking Create Now.

Figure 7-8 The Shadow Copies tab of a volume's Properties dialog box

Caution If you click Disable, you delete all copies that were created by VSS. Consider carefully whether you want to disable VSS for a volume or whether you might be better served by modifying the schedule to prevent new shadow copies from being made.

The default settings configure the server to make copies of shared folders at 7:00 A.M. and noon, Monday through Friday; and 10 percent of the drive space, on the same drive as the shared folder, is used to cache shadow copies.

Each of the following settings can be modified by clicking Settings on the Shadow Copies tab:

- **Storage volume** To enhance performance (not redundancy), you can move the shadow storage to another volume. This must be done when no shadow copies are present. If shadow copies exist, and you want to change the storage volume, you must delete all shadow copies on the volume, then change the storage volume.
- **Details** The dialog box lists shadow copies that are stored and space utilization statistics.
- **Storage limits** This can be as low as 100 MB. When the shadow copy runs out of storage, it deletes older versions of files to make room for newer versions. The proper configuration of this setting depends on the total size of shared folders on a volume with shadowing enabled; the frequency with which files change, and the size of those files; and the number of previous versions you wish to retain. In any event, a maximum of 63 previous versions will be stored for any one file before the earliest version is removed from the shadow storage.

- **Schedule** You can configure a schedule that reflects the work patterns of your users, ensuring that enough previous versions are available without prematurely filling the storage area and thereby forcing the removal of old versions. Remember that when a shadow copy is made, any files that have changed since the previous shadow copy are copied. If a file has been updated several times between shadow copies, those interim versions will not be available.

Using Shadow Copy

Shadow copies of shared folders allow you to access previous versions of files that the server has cached on the configured schedule. This will allow you to

- Recover files that were accidentally deleted
- Recover from accidentally overwriting a file
- Compare versions of files while working

To access previous versions, click the properties of a folder or file and click the Previ ous Versions tab, as shown in Figure 7-9.

Figure 7-9 The Previous Versions tab of a shared resource

The Previous Versions page will not be available if Shadow Copies is not enabled on the server, or if there are no previous versions stored on the server. It will also be unavailable if the shadow copy client has not been installed on your system. This file is located in the %*Systemroot*%\System32\Clients\Twclient\x86 folder of a Windows Server 2003 system. The Windows Installer (.msi) file can be deployed using Group Policy, SMS, or an e-mail message. Finally, the Previous Versions page is only available when accessing a file's properties through a shared folder. If the file is stored on the local hard drive, you will not see the Previous Versions tab, even if the file is shared and VSS is enabled. See this lesson's Practice for an example.

You can then choose to Restore the file to its previous location or Copy the file to a specific location.

Exam Tip Unlike a true restore operation, when you restore a file with Previous Versions, the security settings of the previous version are not restored. If you restore the file to its original location, and the file exists in the original location, the restored previous version overwrites the current version and uses the permissions assigned to the current version. If you copy a previous version to another location, or restore the file to its original location but the file no longer exists in the original location, the restored previous version inherits permissions from the parent folder.

If a file has been deleted, you obviously cannot go to the file's Properties dialog box to locate the Previous Versions page. Instead, open the Properties of the parent folder, click the Previous Versions tab and locate a previous version of the folder that contains the file you want to recover. Click View and a folder window will open, as shown in Figure 7-10, that displays the contents of the folder as of the time at which the shadow copy was made. Right-click the file and choose Copy, then paste it into the folder where you want the file to be recreated.

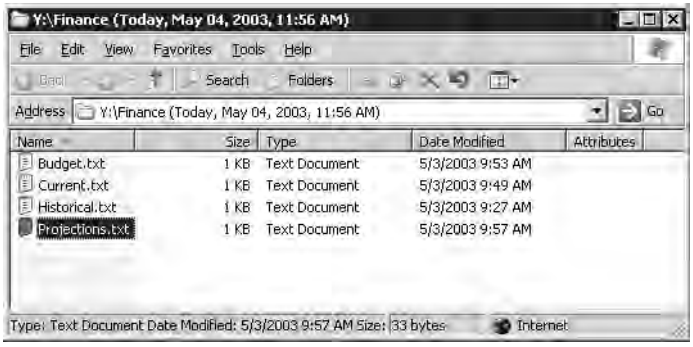

Figure 7-10 A folder's Previous Versions content list

Shadow copy, as you can see, is a useful addition to the toolset for managing file servers and shared data. With VSS, you can preserve data sets at scheduled points in time. Administrators or users can then restore deleted or corrupted files, or compare files to previous versions. As the VSS cache fills, old versions are purged and new shadow copies are added.

If a user requires data to be restored and that data is no longer available through Previous Versions, you can restore the data from backup. If the server becomes corrupted, you must restore the data from backup. Although VSS enhances the manageability and resiliency of shared files, there is no substitute for a carefully planned and verified backup procedure.

Practice: Advanced Backup and Restore

In this practice, you will schedule a backup job, execute a backup from a command prompt, and configure and use Shadow Copies of Shared Folders.

Exercise 1: Schedule a Backup Job

1. Open the Backup Utility and click the Backup tab.
2. From the Job menu, load the Finance Backup selections.
3. Configure the Backup Media Or File Name: C:\Backup-Everyday.bkf.
4. Click Start Backup.
5. Click Advanced and configure an Incremental backup type. Click OK.
6. Click Schedule.
7. In the Set Account Information dialog box, type your password and click OK.
8. Name the job Daily Incremental Backup.
9. Click Properties. Configure the job to run daily. Configure the time to be two min utes from the current time so that you can see the results of the job.
10. Complete configuration of the scheduled job. You will be prompted to enter your password again.
11. Close the Backup Utility.
12. Open the C drive in Windows Explorer and wait two minutes. You will see the backup job appear.
13. Open the Backup Utility, choose the Report command from the Tools menu and view the most recent backup log to confirm the status of the backup job. The num ber of files copied may be zero if you have not made changes to any of the files.
14. If the job did not run properly, open Event Viewer from the Administrative Tools folder. Examine the Application Log to identify the cause of the failure.

Exercise 2: Run a Backup from a Command Prompt

One of the easier ways to determine the correct switches to use for a command prompt backup is to schedule a backup, as you did in Exercise 1, and then examine the com mand that the scheduled task creates.

1. Open the Backup Utility and click the Schedule Jobs tab.
2. Click the icon, in the calendar, representing the scheduled job.
3. Click Properties.
4. Select the command in the Run box and press Ctrl+C to copy it.

5. Cancel to exit the Schedule Jobs dialog box and close the Backup Utility.
6. Open the command prompt.
7. Click the window menu (the icon of the command prompt in the upper-left corner of the command prompt window) and, from the Edit menu, choose Paste. The Ntbackup command with all of its switches is pasted into the command prompt. Press Enter. The backup job is executed.

Note It is recommended that you delete the scheduled backup job at this point in the Practice. You will schedule additional jobs in the Case Scenario, and it will be easier to work with those jobs if the current schedule is clear. In the Backup Utility, click the Schedule Jobs tab, then, in the calendar, click the icon representing the scheduled job. Click Delete.

Exercise 3: Enable Shadow Copies

1. Ensure that the C:\Data folder is shared and that the share permissions are configured to allow Everyone Full Control.
2. Open My Computer.
3. Right-click the C drive and choose Properties.
4. Click the Shadow Copies tab.
5. Select the C volume and click Enable.
6. A message will appear. Click Yes to continue.

Exercise 4: Simulate Changes to Network Files

1. Open the C:\Data\Finance folder and open Current.txt. Modify the file's contents, then save and close the file.
2. Delete the file C:\Data\Finance\Projections.txt.

Exercise 5: Recover Files Using Previous Versions

1. Open the data share by clicking Start, choosing Run, and then typing **\\server01\data**.

Note It is critical that you open the folder using its UNC, not its local path. The Previous Versions tab is only available when connected to a shared folder over the network.

2. Open the Finance folder.
3. Right-click the Current.txt file and choose Properties.

4. Select the Previous Versions tab.
5. Select the previous version of Current.txt.
6. Click Copy, select the Desktop as the destination, and then click Copy again.
7. Click OK to close the Properties dialog box.
8. Open Current.txt from your desktop. You will see that it is the version without the changes you made in Exercise 4.
9. Return to \\Server01\Data. This time, do not open the Finance folder.
10. To recover the deleted Projections.txt file, right-click the Finance folder and click Properties.
11. Select the Previous Versions tab.
12. Select the previous version of the Finance folder and click View.

    A window opens showing the contents of the folder as of the time that the shadow copy was made.
13. Right-click the Projections.txt file and choose Copy.
14. Switch to the folder that shows you the current \\server01\data folder.
15. Open the Finance folder.
16. Paste the Projections.txt file into the folder. You have now restored the previous version of Projections.txt.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Scott Bishop is a power user at a remote site that includes 20 users. The site has a Windows Server 2003 system providing file and print servers. There is a tape drive installed on the system. Because there is no local, full-time administrator at the site, you want to allow Scott to back up and restore the server. However, you want to minimize the power and the privileges that Scott obtains, limiting his capabilities strictly to backup and restore. What is the best practice to provide Scott the minimum necessary credentials to achieve his task?

   _____

   _____

2. Write the command that will allow you to fully back up the C:\Data\Finance folder to a file called Backup.bkf in a share called Backup on Server02, with the backup job name "Backup of Finance Folder." Then, write the command that will allow you to perform an incremental backup and append the backup set to the same file, with the same backup job name.

   _____

   _____

3. A user has deleted a file in a shared folder on a server. The user opens the properties of the folder and does not see a Previous Versions tab. Which of the following may be true? (Choose all that apply.)
   a. The folder is not enabled for Shadow Copy.
   b. The volume on the server is not enabled for Shadow Copy.
   c. The user doesn't have permission to view the Shadow Copy cache.
   d. The Shadow Copy client is not installed on the user's machine.
   e. The folder is on a FAT volume.

Lesson Summary

- You must have the right to backup and restore files to use the Backup Utility or any other backup tool. The right is assigned, by default, to the Backup Operators and Administrators groups.
- The Options dialog box allows you to configure General, Backup, and Restore settings, many of which become defaults that will drive the behavior of the Backup Utility and the Ntbackup command, unless overridden by job-specific options specified in the backup job's Advanced Backup Options dialog box, or in command-line switches.
- The Ntbackup command and its full complement of switches allows you to launch a backup job from a command prompt or batch file.
- Backup jobs can be scheduled to run regularly and automatically during periods of low utilization.
- Volume Shadow Copy Service (VSS) allows a user to access previous versions of files and folders in network shares. With those previous versions, users can restore deleted or damaged files or compare versions of files.

Case Scenario Exercise

You are asked to configure a backup strategy for the Finance Department's shared folder. The backup should occur automatically during the early-morning hours, as there are users working shifts from 4:00 A.M. to 12:00 midnight, Monday through Friday. Files in the folder change frequently—about half the files change once a week; the other half of the files change almost daily. You are told that if the server's hard drive ever fails, down time is extraordinarily costly to the company, so recovery should be as fast as possible.

1. With the knowledge that so many files change almost daily, and that recovery must be as quick as possible, what type of backup job should you consider running nightly?

   Consider normal backups. There is so much change happening to the shared folder, that you are receiving less than a 50 percent benefit using a differential or incremental backup versus a normal backup; and nothing is faster to restore than a normal backup, because the backup set contains all the files to restore.

2. You configure a normal daily backup job to run at 12:00 midnight, after the last shift has gone home for the evening. Unfortunately, you find that the backup job is not completed by 4:00 A.M. when the morning shift arrives. How should you modify your backup strategy?

   Create a normal backup once a week, perhaps on Sunday, and then create differential backups nightly during the week. While differential and incremental backups are both available, differential backups provide faster restore capability, as the most recent differential backup set includes all files that have been updated since the normal backup.

Exercise 1: Create Sample Data

1. Open My Computer and the C drive.
2. Delete the Data folder. You will be prompted to confirm the choice. You will also be informed that the folder is shared, and that deleting the folder will delete the shared folder. Confirm your understanding of the warning and continue.
3. Open the command prompt and type **cd c:**.
4. Type the command **createfiles.bat**.

Note If you did not create the createfiles.bat file in Lesson 1, Exercise 1, complete steps 1 through 3 of Exercise 1 to create the appropriate script.

Exercise 2: Schedule the Backup Job

Configure and schedule the following backup jobs. If you need guidance to achieve these tasks, refer to the instructions in the Practices in Lesson 1 and Lesson 3.

- Normal backup job to back up the C:\Data\Finance folder to a file called C:\BackupFinance.bkf (replacing the media), every Sunday at 9:00 P.M.
- Differential backup job to back up the same folder to the same file (appending to the media), at 12:15 A.M. on Tuesday through Saturday (that is, Monday night through Friday night).

Exercise 3: Simulate the Scheduled Jobs

Rather than waiting until Sunday night for the normal backup job to execute automat ically, you will execute the backup job from the command prompt.

1. Open the Backup Utility.
2. Click the Schedule Jobs tab.
3. Click the icon in the calendar representing the Sunday night normal backup job.
4. Click Properties.
5. Select the command in the Run box and press Ctrl+C to copy it.
6. Cancel to exit the Schedule dialog box and close the Backup Utility.
7. Open the command prompt.
8. Click the window menu (the icon of the command prompt in the upper-left corner of the command prompt window) and, from the Edit menu, choose Paste. The Ntbackup command with all its switches is pasted into the command prompt. Press Enter. The backup job is executed.
9. Open C:\Data\Finance\Projections.txt and make changes to the file. Save and close the file.
10. Repeat steps 1-8, this time executing from the command prompt the *differential* backup job that is scheduled to run every night.

Exercise 4: Verify the Procedure

1. Open the Backup Utility.
2. From the Tools menu, click Report.
3. Open the two most recent backup reports and confirm that the jobs completed successfully. The normal job should have backed up four files. The differential job should have backed up one file.
4. Perform a test restore to a folder called C:\TestRestore. Restore the normal job and then the differential job. If you need guidance, refer to the Practice in Lesson 2.

Caution Remember, before restoring the differential job, that you must configure the Restore options (from the Tools menu, select Options) to always replace files. You may also need to catalog the file to see all the backup sets it contains.

Troubleshooting Lab

At 1:00 P.M. on Tuesday, a user in the Finance Department contacts you to let you know that he accidentally deleted some files from the Finance folder. You are confident that the backup procedure you established will help you recover the deleted files. However, you also want to ensure that you don't roll back any files that had been changed today, after the overnight backup job was executed.

In this lab, you will simulate the workflow that creates such a scenario, and then you will recover the missing data.

Exercise 1: Create a Data Loss

1. Open the C:\Data\Finance folder.
2. Open the file Current.txt. Make some changes to the file. Save and close the file.
3. Open the Budget file. Make some changes, save, and close the file.
4. Delete the Historical.txt and Projections.txt files.

Exercise 2: Plan the Recovery

Review the backup strategy you developed in the Case Scenario Exercise: a normal backup every Sunday night and a differential backup every weeknight.

1. How will you recover the missing data?

   A normal backup includes all selected files. It is the baseline from which you begin to recover from data loss. The differential backup includes all files that have changed since the normal backup. After you have restored the normal backup, you can restore the most recent differential backup. Keep in mind, however, that some of the files (Budget and Current) have been changed by users subsequent to the overnight differential backup.

2. How will you prevent those newer files from being overwritten by files in the backup set?

   The Options dialog box includes a Restore Options tab which allows you to specify how files in the backup set are written to the destination. You can direct the Backup Utility to overwrite files only if the files on the disk are older than the files in the backup set. Files that are newer will remain.

Exercise 3: Recover the Data

1. Open the Backup Utility.
2. Choose the Options command from the Tools menu.
3. Click the Restore tab.
4. Configure restore to leave newer files untouched by selecting Replace The File On Disk Only If The File On Disk Is Older, then close the Options dialog box.
5. Select the backup media that contains your normal and differential backup.
6. Restore the normal backup to its original location.
7. Restore the differential backup to its original location.
8. Open the Current and Budget files. Because these files were newer than those on the backup set, and because of the restore options you configured, they should include the changes you made in the Case Scenario exercise.

Chapter Summary

- You must have the right to back up and restore files to use the Backup Utility or any other backup tool. The right is assigned, by default, to the Backup Operators and Administrators groups.
- The Backup Utility, Ntbackup, allows you to back up and restore data from local and remote folders to local files, tape drives, removable media, or shared folders on remote servers. You cannot back up to writable CD or DVD formats.
- A backup strategy typically begins with a normal backup followed by regular incremental or differential backups. Incremental jobs create the backup more quickly; differential backups are faster to restore. Jobs can be scheduled to occur during periods of low utilization.
- Copy backups and daily backups can be used to capture files without interfering with the regular backup schedule.
- The Backup Utility will also allow you to restore backed up data to the original location or to an alternate location. The latter is useful to test and verify restore procedures. You can control, through the Options dialog box, Restore tab, which files are replaced during a restore.
- The Ntbackup command and its full complement of switches allows you to launch a backup job from a command prompt or batch file.
- Volume Shadow Copy Service (VSS) allows a user to access previous versions of files and folders in network shares. With those previous versions, users can restore deleted or damaged files or compare versions of files.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Identify the group memberships or rights required to perform a backup or restore operation.
- Create a backup strategy based on requirements including the amount of time it takes to back up data, and the speed with which restores must be performed.
- Understand how to restore data under a variety of conditions, including complete and partial data loss. Compare the data loss to the backup schedule to identify the backup sets that must be restored. Integrate your knowledge of the order in which backup sets should be restored and how existing files on the hard drive should be replaced.
- Schedule a backup job and configure backup options.
- Enable shadow copies of shared folders and recover data using the Previous Versions tab of a file or folder's Properties dialog box.

Key Terms

Copy, daily, differential, incremental and normal backup These five backup types select files to back up using specific criteria. *Copy* and *normal* back up all files; *daily* backs up files that have been modified on a specified date; *differential* and *incremental* back up files with their archive attribute set. *Normal* and *incremental* backups also reset the archive attribute.

Archive attribute An attribute that is set when a file is created or modified. Incremental and differential backups will back up files with their Archive attribute set. Incremental backups also clear the Archive attribute.

Volume Shadow Copy Service (VSS) A feature of Windows Server 2003 that allows you to back up files that are locked or open.

Media pools: unrecognized, import, free, backup The four categories of removable media. Ntbackup will back up to media in the free and backup media pools only.

Shadow copies of shared folders A feature of Windows Server 2003 that, once configured on the server and on clients, allows users to retrieve previous versions of files without administrator intervention.

8 Printers

Exam Objectives in this Chapter:

- Monitor print queues.
- Monitor file and print servers. Tools might include Task Manager, Event Viewer, and System Monitor.

Why This Chapter Matters

An administrator's to-do list usually teems with items relating to printers. Whether testing or deploying new printer hardware, troubleshooting print jobs, or securing and monitoring printer utilization, you are apt to be almost as busy with printers as with file and folder access.

Microsoft Windows Server 2003 provides a powerful feature set to support enterprise print services. This chapter introduces you to the setup and configuration of printers on Windows Server 2003, the interaction between printers and the Microsoft Active Directory directory service, connecting clients to network printers, and monitoring and troubleshooting print services. You will learn how to administer local, network, and Internet printers, and how to configure printers for maximum flexibility and security.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to administering Windows Server 2003 printers. This training kit presumes you have a minimum of 18 months of experience and a working knowledge of Active Directory and the Microsoft Management Console (MMC). However, because many administrators come to Windows Server 2003 from other printer environments including Novell NetWare, and because printer terminology has changed slightly, this chapter's first lesson reviews fundamentals of printer configuration. Lesson 2 and Lesson 3 build on those fundamentals to prepare you for advanced, flexible administration, support, monitoring, and troubleshooting, of printers in a Windows Server 2003 environment.

Although it is advantageous to have a printer and two computers (a Windows Server 2003 computer and a client running Windows XP or Windows 2000 Professional), you can complete the exercises in this chapter without a printer, and with only one computer. Prepare the following:

- A Windows Server 2003 (Standard or Enterprise) installed as Server01 and configured as a domain controller in the domain *contoso.com*
- A first-level organizational unit (OU) called Security Groups
- The Active Directory Users And Computers console, or a customized console with the Active Directory Users And Computers snap-in

Lesson 1: Installing and Configuring Printers

Windows Server 2003 supports powerful, secure, and flexible print services. By using a Windows Server 2003 computer to manage printers attached locally to the computer or attached to the network, such printers can be made available to applications running locally on the Windows Server 2003 computer or to users on any client platform, including previous versions of Windows, as well as Netware, UNIX, or Apple Macintosh clients. This lesson will examine the basic concepts, terminology, and skills related to the setup of printers in Windows Server 2003.

After this lesson, you will be able to

- Understand the model and terminology used for Windows printing
- Install a logical printer on a print server for a network attached printer
- Prepare a print server to host clients including computers running previous versions of Windows
- Connect a printer client to a logical printer on a print server
- Manage print jobs

Estimated lesson time: 15 minutes

Understanding the Windows Server 2003 Printer Model

Windows Server 2003, and previous versions of Windows, support two types of printers:

- **Locally attached printers** Printers that are connected to a physical port on a print server, typically a universal serial bus (USB) or parallel port.
- **Network-attached printers** Printers connected to the network instead of a physical port. A network-attached printer is a node on the network; print servers can address the printer using a network protocol such as Transmission Control Protocol/Internet Protocol (TCP/IP).

Each type of printer is represented on the print server as a logical printer. The *logical printer* defines the characteristics and behavior of the printer. It contains the driver, printer settings, print setting defaults and other properties that control the manner in which a print job is processed and sent to the chosen printer. This virtualization of the printer by a logical printer allows you to exercise extraordinary creativity and flexibility in configuring your print services.

Note In previous versions of Windows and in earlier versions of documentation, the printer was referred to as the "print device" and the logical printer was referred to as the "printer."

There are two ways to implement printing to network attached printers. One model is created by installing logical printers on all computers, and connecting those logical printers directly to the network-attached printer. In this model, there is no print server; each computer maintains its own settings, print processor, and queue. When users examine the print queue, they see only the jobs they have sent to the printer. There is no way for users to know what jobs have been sent to the printer by other users. In addition, error messages appear only on the computer that is printing the current job. Finally, all print job processing is performed locally on the user's computer, rather than being offloaded to a print server.

Because of these significant drawbacks, the most typical configuration of printers in an enterprise is a three-part model consisting of the physical printer itself, a logical printer hosted on a print server, and printer clients connecting to the server's logical printer. This lesson focuses exclusively on such a structure, although the concepts and skills discussed apply to other printer configurations.

Printing with a print server provides the following advantages:

- The logical printer on the print server defines the printer settings and manages printer drivers.
- The logical printer produces a single print queue that appears on all client computers, so users can see where their jobs are in relation to other users' jobs.
- Error messages, such as out-of-paper or printer-jam messages, are visible on all clients, so all users can know the state of the printer.
- Most applications and most print drivers will offload some, or a significant amount, of the print-job processing to the server, which increases the responsiveness of the client computers. In other words, when users click Print, their jobs are sent quickly to the print server and users can resume their work while the print server processes the jobs.
- Security, auditing, monitoring, and logging functions are centralized.

Installing a Printer on Windows Server 2003

Printers are managed most commonly through the Printers And Faxes folder, which integrates both printer and fax capabilities. The Add Printer Wizard guides you through the printer setup. The most critical choices you must make are the following:

- **Local Or Network Printer** This page of the Add Printer Wizard is shown in Figure 8-1. When you set up a printer on a Windows Server 2003 computer, the terms local printer and network printer have slightly different meanings from what you might expect. A *local printer* is a logical printer that supports a printer attached directly to the server or a stand-alone, network-attached printer. When you direct the Add Printer Wizard to create a local printer by clicking Local Printer Attached

To This Computer, the server can share the printer to other clients on the network. A *network printer*, on the other hand, is a logical printer that that connects to a printer directly attached to another computer or to a printer managed by another print server. The user interface can be misleading, so remember that, in the common print server implementation, the print server will host local printers (whether the printer hardware is attached to the computer or network-attached), and workstations will create network printers connecting to the server's shared logical printer.

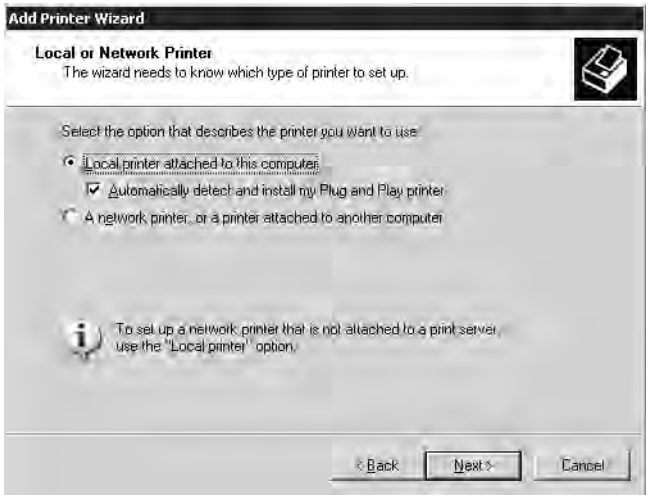

Figure 8-1 The Local Or Network Printer page of the Add Printer Wizard

- **Select A Printer Port** When you create a local printer on a print server, the Add Printer Wizard asks you to specify the port to which the printer is attached. If the port already exists, whether a local port such as LPT1 or a network port specified by an IP address, select the port from the Use The Following Port drop-down list. When setting up a logical printer for a network attached printer for which a port has not been created, click Create A New Port, select Standard TCP/IP Port and click Next. The Add Standard TCP/IP Printer Port Wizard appears. Clicking Next prompts you for the IP address or DNS name of the printer. After the port has been added, you are returned to the Add Printer Wizard.
- **Install Printer Software** If Plug and Play does not detect and install the correct printer automatically, you can select your printer from an extensive list that is categorized by manufacturer. If the printer does not appear on the list, you can click Have Disk and install the printer from drivers supplied by the manufacturer.
- **Printer Name and Share Name** Although Windows Server 2003 supports long printer names and share names including spaces and special characters, it is best practice to keep names short and simple. The entire qualified name including the server name (for example, \\Server01\PSCRIPT) should be 32 characters or fewer.

The share name and the printer name appear, and are used in different places throughout the Windows user interface. Although the share name is independent of, and can be different from, the printer name, many enterprises unify the printer name and the share name to reduce confusion.

Configuring Printer Properties

After installing the logical printer, you can configure numerous properties by opening the printer's Properties dialog box, shown in Figure 8-2. The General tab allows you to configure the printer name, location, and comments, all of which were initially configured based on your responses to prompts in the Add Printer Wizard.

Figure 8-2 The General tab of a printer's Properties dialog box

The Sharing tab shown in Figure 8-3 allows you to specify whether the logical printer is shared, and is therefore available to other clients on the network, and whether the printer is listed in Active Directory, a default setting, for shared printers, that allows users to easily search for and connect to printers.

Note You can use the Sharing tab to stop sharing a printer, if you take a printer offline and want to prevent users from accessing the printer.

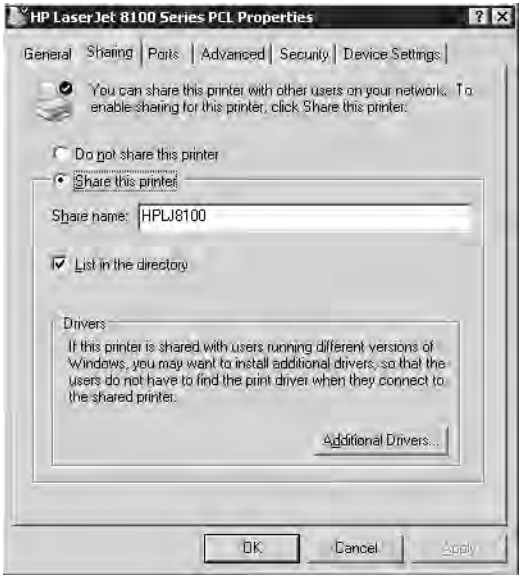

Figure 8-3 The Sharing tab of a printer's Properties dialog box

During printer setup, Windows Server 2003 loads drivers onto the print server that support that printer for clients running Windows Server 2003, Windows XP, and Windows 2000. Printer drivers are platform-specific. If other platforms will be connecting to the shared logical printer, install the appropriate drivers on the server, so that Windows clients will download the driver automatically when they connect. Otherwise, you will be prompted for the correct drivers on each individual client.

On the Sharing tab of the Properties dialog box, click Additional Drivers to configure the print server to host drivers for computers running versions of Windows prior to Windows 2000. When you select a previous version of Windows, the server will prompt you for the drivers for the appropriate platform and printer. Those drivers will be available from the printer's manufacturer, or sometimes on the original CD-ROM of the previous version of Windows.

By loading drivers on the server for all client platforms, you can centralize and facilitate driver distribution. Client computers running Windows NT, Windows 2000, Windows XP, and Windows Server 2003 download the driver when they first connect to the shared printer. They also verify that they have the current printer driver each time they print and, if they do not, they download the updated driver. For these client computers, you need only update printer drivers on the print server. Client computers running Windows 95 or Windows 98 do not check for updated printer drivers, once the driver is initially downloaded and installed. You must manually install updated printer drivers on these clients.

Other printer properties will be discussed later in this chapter.

Tip You can access other servers' printer folders by browsing the network or by choosing the Run command from the Start menu and typing **\\server_name**. You can drag those servers' Printer and Faxes folders to your own, giving you easy access to manage remote printers.

Connecting Clients to Printers

Printers that have been set up as logical printers on a print server can be shared to other systems on the network. Those systems will also require logical printers to represent the network printer.

Configuring a print client can be done in several ways, including the Add Printer Wizard, which can be started from the Printers And Faxes folder or from the common Windows Print dialog box in almost all Microsoft applications, including Internet Explorer and Notepad. On the Local or Network Printer page, select A Network Printer Or A Printer Attached To Another Computer. When prompted for the printer name, you can search Active Directory, enter the Universal Naming Convention (UNC) (for example, \\Server\Printersharename) or Uniform Resource Locator (URL) to the printer, or browse for the printer using the Browser service.

One of the more efficient ways to set up print clients is to search Active Directory for the printer. In the Specify A Printer page of the Add Printer Wizard, choose Find A Printer In The Directory and click Next. The Find Printers dialog box appears, as shown in Figure 8-4, and you can enter search criteria including printer name, location, model, and features. Wildcards can be used in many of the criteria. Click Find Now and a result set is displayed. Select the printer and click OK. The Add Printer Wizard then steps you through remaining configuration options.

Tip You can save a search by choosing Save Search from the File menu. As an administrator, you can create and save custom searches to users' desktops, allowing them to easily locate predefined subsets for the printers in your enterprise.

A logical printer includes the drivers, settings, and print queue for the printer on the selected port. When you double-click a printer in the Printers And Faxes folder, a window opens that displays the jobs in the printer's queue. By right-clicking any job, you can pause, resume, cancel, or restart the job. From the Printer menu, you can also pause or cancel all printing, access the printer properties, or set the printer as default or offline. Your ability to perform each of these actions depends, of course, upon the permissions on the printer's access control list.

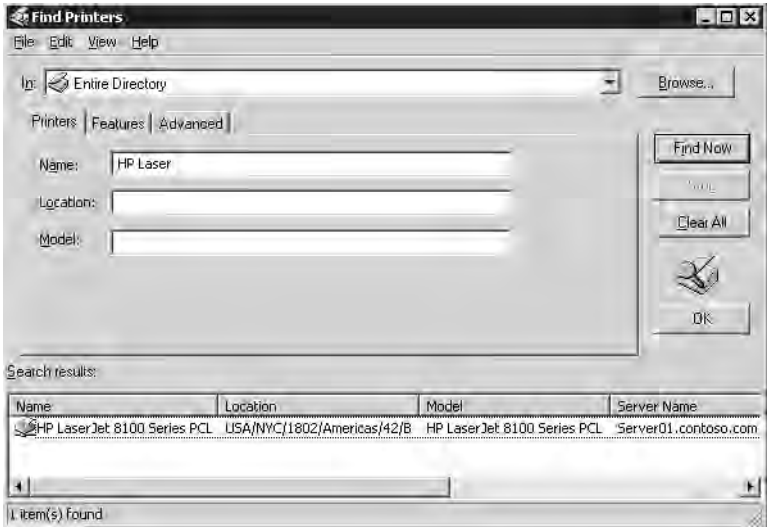

Figure 8-4 The Find Printers dialog box

As an alternative to using the Add Printer Wizard, if you are using Windows Server 2003 or Windows XP with the default Start menu, perform the following steps to configure a print client:

1. Click Start, and then select Search.
2. In the Search Companion pane, click Other Search Options, then Printers, Computers, Or People, and finally A Printer On The Network.
3. The Find Printers dialog box will be displayed, allowing you to search for the printer using various criteria.
4. After entering the desired criteria, click Find Now.

Practice: Installing and Configuring a Printer

In this practice, you will set up a logical printer on a print server and simulate connecting a client to the shared printer. You will then send a print job to the printer.

You do not need to have a print device connected to Server01 or to the network, nor are you required to have a second computer to act as a print client. However, if you have access to these additional components, you are encouraged to implement the exercises using that extra hardware.

Exercise 1: Add a Local Printer and Configure Print Sharing

In this exercise, you use the Add Printer Wizard to add a logical printer to Server01. The printer will connect to a network-attached HP LaserJet 8100 that is connected to

the network at IP address 10.0.0.51. You do not need an actual printer to complete this exercise.

1. Log on to Server01 as Administrator.
2. Open the Printers And Faxes folder.
3. Double-click Add Printer. The Add Printer Wizard appears.
4. Click Next. The Local Or Network Printer page appears.

   You are prompted for the location of the printer. Although the printer is attached to the network, the logical printer serving that printer is being added to Server01, so the printer is referred to as a local printer.
5. Verify that the Local Printer option is selected and that the Automatically Detect And Install My Plug And Play Printer check box is cleared (because you are configuring a printer for a fictional device), and then click Next.
6. The Select A Printer Port page appears. Click Create A New Port.
7. Select Standard TCP/IP Port from the Type Of Port drop-down list.

   The port types that will be available, other than local port, depend on the installed network protocols. In this case, TCP/IP is installed, so this protocol-based port is available.
8. Click Next. The Add Standard TCP/IP Printer Port Wizard appears.
9. Click Next.
10. Enter the IP Address: 10.0.0.51 and accept the default port name, IP_10.0.0.51.
11. Click Next.

    Because a print device is not actually attached to the network at that address, there will be a delay while the Wizard attempts to locate and identify the printer. You will also be prompted to specify the type of network interface.
12. Select Hewlett Packard Jet Direct as the device type.
13. Click Next, and then click Finish. The Add Standard TCP/IP Printer Port Wizard closes, returning you to the Add Printer Wizard.

    The Wizard prompts you for the printer manufacturer and model. You will add an HP LaserJet 8100 Series PCL printer.

Tip The printers list is sorted in alphabetical order. If you cannot find a printer name, make sure that you are looking in the correct location.

14. From the Manufacturer list, click HP; from the Printers list, scroll down the list, click HP LaserJet 8100 Series PCL; and then click Next.

    The Name Your Printer page appears. The default name in the Printer Name field is the printer model, HP LaserJet 8100 Series PCL. The name you enter should conform to naming conventions in your enterprise. For this exercise, enter the name HPLJ8100.

15. Type **HPLJ8100** and Click Next.

    The Printer Sharing page appears, prompting you for printer-sharing information. The share name should also reflect naming conventions in your enterprise. As discussed earlier, the printer's UNC (that is, \\Servername\Printersharename) should not exceed 32 characters.

16. Verify that the Share Name option is selected.

17. In the Share Name text box, type **HPLJ8100**, and then click Next.

    The Location And Comment page appears.

Note The Add Printer Wizard displays the values you enter for the Location and Comment text boxes when a user searches the Active Directory for a printer. Entering this information is optional, but doing so helps users locate the printer.

18. In the Location text box, type **USA/NYC/1802Americas/42/B**.

19. In the Comment text box, type **Black and White Output Laser Printer-High Volume**.

20. Click Next.

    The Print Test Page screen appears. A test page that prints successfully would confirm that your printer is set up properly.

21. Choose No (because the printer doesn't exist) and click Next. The Completing The Add Printer Wizard page appears and summarizes your installation choices.

22. Confirm the summary of your installation choices, and then click Finish.

    An icon for the printer appears in the Printers And Faxes window. Notice that Windows Server 2003 displays an open hand beneath the printer icon. This indicates the printer is shared. Also notice the check mark next to the printer, which indicates the printer is the default printer for the print server.

23. Keep the Printers And Faxes window open because you will need it to complete the next exercise.

Exercise 2: Connect a Client to a Printer

If you have access to a second computer, you would install on each workstation a printer that connects to the shared printer on Server01. In this practice, you are required to have only one computer (Server01), but you can simulate connecting a printer client to the server's logical printer.

1. Open the Printers And Faxes folder.
2. Start the Add Printer Wizard and click Next.
3. In the Local Or Network Printer dialog box, select A Network Printer, Or A Printer Attached To Another Computer and click Next.
4. Confirm that Find A Printer In The Directory is selected and click Next. The Find Printers dialog box appears.
5. In the Location box, type ***NYC*** and then click Find Now.
6. Select the printer HPLJ8100 in the results list and click OK.
7. On the Add Printer Wizard's Default Printer page, select Yes and then click Next.
8. Click Finish.

   You will *not* see a new printer icon in the Printers And Faxes folder because it is not possible to create a printer client to a logical printer on the same computer. If you conduct this exercise on a second computer, you will see the icon for the new printer appear.

Exercise 3: Take a Printer Offline and Print a Test Document

In this exercise, you set the printer you created to offline status. Taking a printer offline causes documents you send to this printer to be held in the print queue while the print device is unavailable. Doing this will prevent error messages about unavailable print devices from occurring in later exercises. Otherwise, Windows Server 2003 will display error messages when it attempts to send documents to the fictional print device that is not actually available to the computer.

1. In the Printers And Faxes window, right-click the HPLJ8100 icon.
2. Choose Use Printer Offline. Notice that the icon appears dimmed to reflect that the printer is not available, and the status appears as Offline.
3. Double-click the HPLJ8100 icon. Notice that the list of documents to be sent to the print device is empty.
4. Click the Start menu, point to Programs, point to Accessories, and then click Notepad.
5. In Notepad, type any sample text that you want.
6. Arrange Notepad and the HPLJ8100 window so that you can see the contents of each.

7. From the File menu in Notepad, select Print. The Print dialog box appears, allowing you to select the printer and print options.

   The Print dialog box displays the location and comment information you entered when you created the printer, and it shows HPLJ8100 as the default and selected printer, and indicates that the printer is offline.

8. Click Print. Notepad briefly displays a message stating that the document is printing on your computer. On a fast computer, you might not see this message.

   In the HPLJ8100–Use Printer Offline window, you will see the document waiting to be sent to the print device. The document is held in the print queue because you took the printer offline. If the printer were online, the document would be sent to the print device.

9. Close Notepad, and click No when prompted to save changes to your document.

10. Select the document in the HPLJ8100 window and, from the Printer menu, select Cancel All Documents. A Printers message box appears, asking if you are sure you want to cancel all documents for HPLJ8100.

11. Click Yes. The document is removed.

12. Close the HPLJ8100–Use Printer Offline window.

13. Close the Printers And Faxes window.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're setting up a printer on your Windows Server 2003 computer. The computer will be used as a print server on your network. You plan to use a print device that's currently connected to the network as a stand-alone print device. Which type of printer should you add to the print server? (Choose all that apply.)

   a. Network

   b. Shared

   c. Local

   d. Remote

2. You're installing a printer on a client computer. The printer will connect to a logical printer installed on a Windows Server 2003 print server. What type or types of information could you provide to set up the printer? (Choose all that apply.)
   a. TCP/IP printer port
   b. Model of the print device
   c. URL to printer on print server
   d. UNC path to print share
   e. Printer driver
3. One of your printers is not working properly, and you want to prevent users from sending print jobs to the logical printer serving that device. What should you do?
   a. Stop sharing the printer
   b. Remove the printer from the Active Directory
   c. Change the printer port
   d. Rename the share
4. You're administering a Windows Server 2003 computer configured as a print server. You want to perform maintenance on a print device connected to the print server. There are several documents in the print queue. You want to prevent the documents from being printed to the printer, but you don't want users to have to resubmit the documents to the printer. What is the best way to do this?
   a. Open the printer's Properties dialog box, select the Sharing tab, and then select the Do Not Share This Printer option.
   b. Open the printer's Properties dialog box and select a port that is not associated with a print device.
   c. Open the printer's queue window, select the first document, and then select Pause from the Document window. Repeat the process for each document.
   d. Open the printer's queue window, and select the Pause Printing option from the Printer menu.

Lesson Summary

- A printer client submits a print job to a print server, which in turn sends the job to the printer. The printer client and the print server each maintain a logical printer representing the printer.
- A local printer is one that supports a printer directly attached to the computer or attached to the network.
- A network printer connects to a logical printer maintained by another computer: a print server.
- Microsoft Windows clients will download the printer driver automatically from the logical printer on the print server. Printers can be added using the printer's Sharing property page.

Lesson 2: Advanced Printer Configuration and Management

In the previous lesson, you learned that the Windows printer model is best leveraged when a logical printer is created to support a physical device—either directly attached to the computer or attached to the network—and when that logical printer is shared to printer clients. That logical printer on the print server becomes a central point of configuration and management. The drivers that you install on the printer are downloaded automatically by Windows clients, and the settings you configure for the printer are distributed as the settings for each of the printer's clients.

This lesson takes this virtualization of printers as logical devices to the next level. After examining printer properties, including printer security, you will learn how to create printer pools to provide faster turnaround for client print jobs. You will also learn how to make better use of your printers by creating more than one logical printer for a device to configure, manage, or monitor print jobs or printer usage more effectively. Finally, you will learn how to manage Active Directory printer objects and Internet printing.

After this lesson, you will be able to

- Manage and configure printer properties
- Create a printer pool
- Configure multiple logical printers to support a single printer
- Manage and connect to printers using Active Directory and Internet Printing Protocol (IPP)

Estimated lesson time: 30 minutes

Managing Printer Properties

Printers and print jobs are managed from their properties dialog boxes. These properties dialog boxes can be accessed from the Printers And Faxes folder. Right-click a printer and select Properties to configure a printer. Double-click a printer and, in the print queue, right-click a print job and choose Properties to configure a print job. The initial properties of a print job are inherited from the properties of the printer itself. But a print job's default properties can be modified independently of the printer's.

Controlling Printer Security

Windows Server 2003 allows you to control printer usage and administration by assigning permissions through the Security tab of the printer's Properties dialog box. You can assign permissions to control who can use a printer and who can administer the printer or documents processed by the printer. A typical printer Security tab of a printer's Properties dialog box is shown in Figure 8-5.

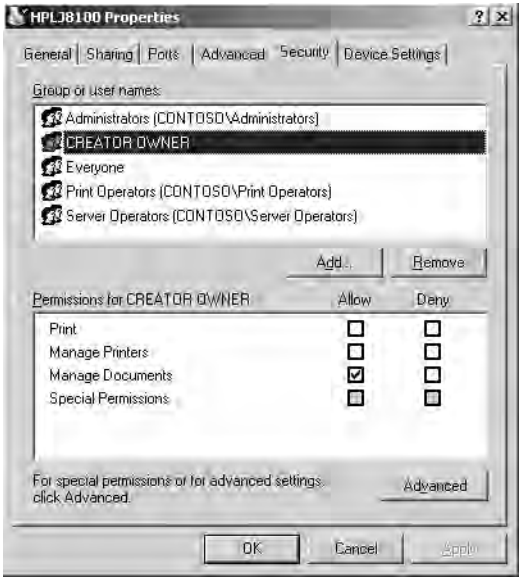

Figure 8-5 The Security tab of a printer's Properties dialog box

You can use a printer's access control list (ACL) to restrict usage of a printer and to delegate administration of a printer to users who are not otherwise administrators. Windows Server 2003 provides three levels of printer permissions: Print, Manage Printers, and Manage Documents.

By default, the Print permission is assigned to the Everyone group. Choosing this permission allows all users to send documents to the printer. To restrict printer usage, remove this permission and assign Allow Print permission to other groups or individual users. Alternatively, you can deny Print permission to groups or users. As with file system ACLs, denied permissions override allowed permissions. Also, like file system ACLs, it is best practice to restrict access by assigning allow permissions to a more restricted group of users rather than granting permissions to a broader group and then having to manage access by assigning additional deny permissions.

The Manage Documents permission provides the ability to cancel, pause, resume, or restart a print job. The Creator Owner group is allowed Manage Documents permission. Because a permission assigned to Creator Owner is inherited by the user that creates an object, this permission enables a user to cancel, pause, resume, or restart a print job that he or she has created. The Administrators, Print Operators and Server Operators groups are also allowed the Manage Documents permission, which means they can cancel, pause, resume, or restart *any* document in the print queue. Those three groups are also assigned the Allow Manage Printers permission, which enables them to modify printer settings and configuration, including the ACL itself.

Tip If a printer's security is not a major concern, you can delegate administration of the printer by assigning a group, such as the <*Printer*> Users group, Manage Documents, or even Manage Printers permission.

Assigning Forms to Paper Trays

If a print device has multiple trays that regularly hold different paper sizes, you can assign a form to a specific tray. A form defines a paper size. When users print a document of a particular paper size, Windows Server 2003 automatically routes the print job to the paper tray that holds the correct form. Examples of forms include Legal, Letter, A4, Envelope, and Executive.

To assign a form to a paper tray, select the Device Settings tab of the printer's Properties dialog box, as shown in Figure 8-6. The number of trays shown in the Form To Tray Assignment section obviously depends on the type of printer you have installed, and the number of trays it supports. Further down the Device Settings tree are settings to indicate the installation state of printer options, such as additional paper trays, paper handling units, fonts, and printer memory.

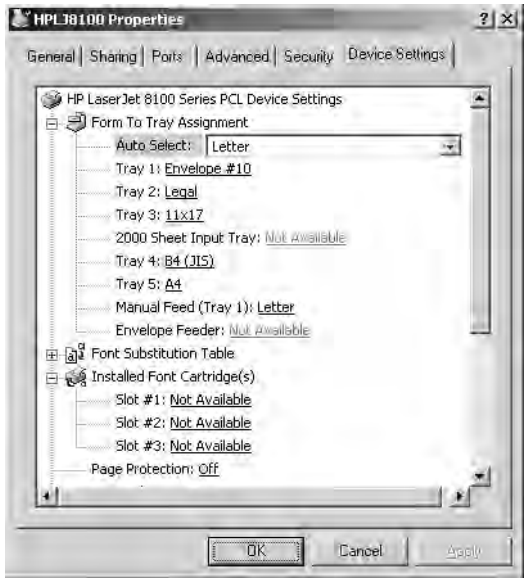

Figure 8-6 The Device Settings tab of a printer's Properties dialog box

Print Job Defaults

The General tab of the printer's Properties dialog box includes a Printing Preferences button, and the Advanced tab includes a Printing Defaults button. Both of these buttons display a dialog box that lets you control the manner in which jobs are printed by the logical printer, including page orientation (portrait or landscape), double-sided

printing (if supported), paper source, resolution, and other document settings. These dialog boxes are identical to each other, and are also identical to the dialog box a user receives when clicking Properties in a Print dialog box.

Why are there three print job Properties dialog boxes? The Printing Defaults dialog box configures default settings for all users of the logical printer. If the printer is shared, its printing defaults become the default properties for all printers connected from clients to the shared printer. The Printing Preferences dialog box configures the user-specific, personal preferences for a printer. Any settings in the Printing Preferences dialog box override printing defaults. The Properties dialog box that can be accessed by clicking Properties in a Print dialog box configures the properties for the specific job that is printed. Those properties will override both printing defaults and printing preferences. This triad of print job property sets allows administrators to configure a printer centrally, by setting printing defaults on the shared logical printer, and allows flexibility and decentralized configuration by users or on a document-by-document basis.

Printer Schedule

The Advanced tab of a printer's Properties dialog box, as shown in Figure 8-7, allows you to configure numerous additional settings that drive the behavior of the logical printer, its print processor and spool. Among the more useful and interesting setting is printer's schedule.

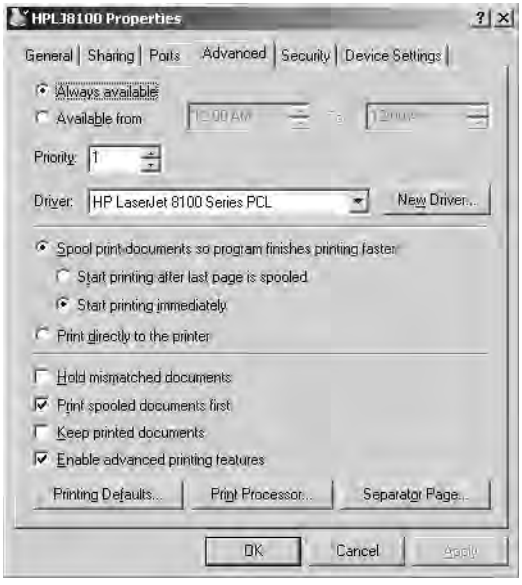

Figure 8-7 The Advanced tab of a printer's Properties dialog box

The logical printer's schedule determines when a job is released from the spool, or queue, and sent to the printer itself. A user with Allow Print permission can send a job to the printer at any time, but the job will be held until the printer's schedule allows it to be directed to the printer's port. Such a configuration is not appropriate for normal, day-to-day printers. However a schedule is invaluable for situations in which users are printing large jobs, and you want those jobs to print after hours, or during periods of low use. By configuring a printer's schedule to be available during night hours, users can send the job to the printer during the day, the printer will complete the jobs overnight, and the users can pick up those printing jobs the next morning.

Tip When you set up a printer pool, place the print devices in the same physical location so that users can easily locate their documents. When users print to a printer pool, there is no way to know which individual printer actually printed the job.

Setting Up a Printer Pool

A printer pool is one logical printer that supports multiple physical printers, either attached to the server, attached to the network, or a combination thereof. When you create a printer pool, users' documents are sent to the first available printer—the logical printer representing the pool automatically checks for an available port.

Printer pooling is configured from the Ports tab of the printer's Properties dialog box. To set up printer pooling, select the Enable Printer Pooling check box, and then select or add the ports containing print devices that will be part of the pool. Figure 8-8 shows a printer pool connected to three network-attached printers.

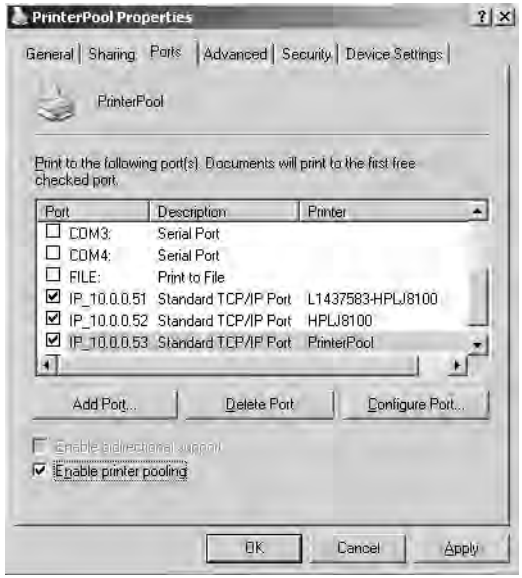

Figure 8-8 The Ports tab of a printer pool's Properties dialog box, showing a three-printer pool

Exam Tip The driver used by the printer pool must be compatible with all printers to which the pool directs print jobs.

Configuring Multiple Logical Printers for a Single Printer

Although a printer pool is a single logical printer that supports multiple ports, or printers, the reverse structure is more common and more powerful: multiple logical printers supporting a single port, or printer. By creating more than one logical printer directing jobs to the same physical printer, you can configure different properties, printing defaults, security settings, auditing, and monitoring for each logical printer.

For example, you might want to allow executives at Contoso Ltd. to print jobs immediately, bypassing documents that are being printed by other users. To do so, you can create a second logical printer directing to the same port (the same physical printer) as the other users, but with a higher priority.

Use the Add Printer Wizard to generate an additional logical printer. To achieve a multiple logical printer-single port structure, additional printers use the same port as an existing logical printer. The printer name and share name are unique. After the new printer has been added, open its properties and configure the drivers, ACL, printing defaults, and other settings of the new logical printer.

To configure high priority for the new logical printer, click the Advanced tab and set the priority, in the range of 1 (lowest) to 99 (highest). Assuming that you assigned 99 to the executives' logical printer, and 1 to the printer used by all users, documents sent to the executives' printer will print before documents queued in the users' printer. An executive's document will not interrupt a user's print job. However, when the printer is free, it will accept jobs from the higher-priority printer before accepting jobs from the lower-priority printer. To prevent users from printing to the executives' printer, configure its ACL and remove the print permission assigned to the Everyone group, and instead allow only the executives' security group print permission.

Exam Tip Remember that a printer pool is a single logical printer serving multiple ports; and all other variations on the standard print client—print server—printer structure are achieved by creating multiple logical printers serving a single port.

Windows Server 2003 Printer Integration with Active Directory

The print subsystem of Windows Server 2003 is tightly integrated with Active Directory, making it easy for users and administrators to search for and connect to printers throughout an enterprise. All required interaction between printers and Active Directory is configured, by default, to work without administrative intervention. You only need to make changes if the default behavior is not acceptable.

When a logical printer is added to a Windows Server 2003 print server, the printer is automatically published to Active Directory. The print server creates a printQueue object and populates its properties based on the driver and settings of the logical printer.

Off the Record The printer objects are not easy to find in Active Directory Users and Computers. You must use the Find Objects In Active Directory button on the MMC toolbar or select View Users, Groups, And Computers As Containers from the View menu, at which point printer objects will become visible inside the print server. The printer is placed in the print server's computer object in the Active Directory service. The object can be moved to any OU.

When any change occurs in the printer's configuration, the Active Directory printer object is updated. All the configuration information is sent again to the Active Directory store even if some of it has remained unchanged.

Planning Creation and updating of printer objects happens relatively quickly, but objects and attributes must be replicated before they affect the results of a Find Printers operation from a client. Replication latency depends on the size of your enterprise, and your replication topology.

If a print server disappears from the network, its printer object is removed from the Active Directory. The printer Pruner service confirms the existence of shared printers represented in Active Directory by contacting the shared printer every eight hours. A printer object will be pruned if the service is unable to contact the printer two times in a row. This might occur if a print server is taken offline. It will happen regularly if printers are shared on Windows 2000 or Windows XP workstations that are shut off overnight or on weekends. However, a print server will recreate the printer objects for its printers when the machine starts, or when the spooler service is restarted. So, again, administrative intervention is not required.

Publishing Windows Printers

Printers that are added by using the Add Printer Wizard are published by default. The Add Printer Wizard does not allow you to prevent the printer from being published to the Active Directory service when you install or add a printer.

If you want to re-publish a printer (for example, after updating its name or other properties), or if you do *not* want a shared printer published in Active Directory, open the printer's Properties dialog box, click the Sharing tab, and select or clear the List In The Directory check box.

Note A printer connected to a local port is likely to be detected and installed automatically by Plug And Play. In this case, you must share and publish the printer manually using the Sharing tab.

Logical printers that are shared on computers running Windows NT 4 or Windows NT 3.51 are not published automatically, but can be manually published using the Active Directory Users And Computers MMC console. Simply right-click the OU or other container in which you want to create the printer and choose New Printer.

Planning You should add only printer objects that map to printers on pre–Windows 2000 computers. Do not add printer objects for printers on computers running Windows 2000 or later; allow those printers to publish themselves automatically.

Manually Configuring Printer Publishing Behavior

All the default system behaviors described above can be modified using local or group policy. Printer policies are located in the Computer Configuration node, under Administrative Templates. For a description of each of these policies, open the Properties dialog box for a specific policy and click the Explain tab.

Printer Location Tracking

Printer location tracking is a feature, disabled by default, that significantly eases a user's search for a printer in a large enterprise by pre-populating the Location box of the Find Printers dialog box, so that the result set will automatically be filtered to list printers in geographic proximity to the user.

To prepare for printer location tracking, you must have one or more sites *or* one or more subnets. Site and subnet objects are created and maintained using the Active Directory Sites And Services MMC snap-in or console. You must also configure the Location tab of the site or subnet Properties dialog box using a naming convention that creates a hierarchy of locations, separated by slashes. For example, the location USA/NYC/1802Americas/42/B might refer to a building at 1802 Avenue of the Americas in Manhattan, on the 42nd floor in Area B. A location may span more than one subnet, or more than one site.

You must then enable printer location tracking using the Pre-Populate Printer Search Location Text policy.

Active Directory is able to identify a computer's site or subnet affiliation based on the computer's IP address. When the Find Printers dialog box is invoked, the computer's location, as defined in its corresponding site or subnet object, will be automatically

placed in the Location box. A Browse button will also appear, enabling a user to browse the location hierarchy for printers in other locations.

This powerful feature simplifies printer administration and setup considerably. However, it obviously requires careful planning on the back end to ensure that all subnets are defined, and that a reasonable, hierarchical location naming convention has been applied consistently. More information about this feature is available in the online Help and Support Center.

Internet Printing

Windows Server 2003 supports an additional set of functionality through the Internet Printing Protocol (IPP), which enables users to connect to printers and send print jobs over encapsulated Hypertext Transfer Protocol (HTTP). Internet printing also gives administrators the option to manage and configure printers using any variety of Internet browsers and platforms.

Setting Up Internet Printing

Internet printing is not installed or enabled by default in Windows Server 2003. You must install Internet Information Services (IIS), as discussed in Chapter 6. Internet printing is available for installation when you install IIS. To install Internet printing, perform the following steps:

1. Open Add/Remove Programs in Control Panel and click Add/Remove Windows Components.
2. Select Application Server and click Details.
3. Select Internet Information Services (IIS) and click Details.
4. Select Internet Printing.

Once IIS and Internet printing are installed, you can disable or enable the feature using the IIS snap-in or console. Expand the server's node and click Web Service Extensions. In the details pane, select Internet Printing, and click Prohibit or Allow.

Internet printing creates a Printers virtual directory under the Default Web site. This virtual directory points to %*Systemroot*%\Web\Printers. The printer site is accessed using Microsoft Internet Explorer 4.01 and later by typing the address of the print server in the Address box followed by the Printers virtual directory name. For example, to access the Internet printing page for Server01, type **http://Server01/printers/**.

Note You can configure authentication and access security for Internet printing using the virtual directory's Properties dialog box.

Using and Managing Internet Printers

You can connect to *http://printserver/printers* to view all printers on the print server. After locating the desired printer and clicking it, a Web page for that printer is displayed.

As a shortcut, if you know the exact name of the printer to which you want to connect, type the address of the printer using the following format:

http://printserver/printersharename/

Once the printer's Web page is displayed, you can connect to or manage the printer, assuming you have been allowed appropriate security permissions. When you click Connect on the printer's Web page, the server generates a .cab file that contains the appropriate printer driver files and downloads the .cab file to the client computer. The printer that is installed is displayed in the Printers folder on the client. The printer can then be used and managed from the Printers And Faxes folder like any other printer. Using a Web browser to manage printers has several advantages:

- It allows you to administer printers from any computer running a Web browser, regardless of whether the computer is running Windows Server 2003 or has the correct printer drivers installed.
- It allows you to customize the interface. For example, you can create your own Web page containing a floor plan with the locations of the printers and the links to the printers.
- It provides a summary page listing the status of all printers on a print server.
- Internet printing can report real-time print device data, such as whether the print device is in power-saving mode, if the printer driver makes such information available. This information is not available from the Printers And Faxes window.

Practice: Advanced Printer Configuration and Management

In this practice, you will configure printer pooling and configure a second logical printer to a single network-attached printer.

Exercise 1: Configure Printer Pooling

1. From the Printers And Faxes window, create a new printer. If you need guidance for how to create a printer, follow the steps in Lesson 1, Exercise 1. The printer should direct to the network address 10.0.0.52 (a new port). Configure the printer as an HP LaserJet 8100 Series PCL, and use PrinterPool as the printer name and the share name. All other properties, including location and comment, are the same as in Lesson 1, Exercise 1.
2. Open the properties of PrinterPool.

3. Click the Ports tab.
4. Select the Enable Printer Pooling check box, and then click the check box next to the port IP_10.0.0.51.
5. Click Apply. Both network ports are now selected.

   Will users sending print jobs to HPLJ8100 benefit from printer pooling?

   No. Printer pooling was configured for the shared printer named PrinterPool. Print jobs sent to PrinterPool can print to the printers at 10.0.0.51 and 10.0.0.52. Print jobs sent to HPLJ8100 can print only to the printer at 10.0.0.51.

Exercise 2: Configure Multiple Logical Printers for a Single Printer

1. From the Printers And Faxes window, create a new printer. If you need guidance for how to create a printer, follow the steps in Lesson 1, Exercise 1. The printer should direct to the network IP address 10.0.0.52 (note the port already exists). Configure the printer as an HP LaserJet 8100 Series PCL, and use PriorityPrinter as the printer name and the share name. All other properties, including location and comment, are the same as in Lesson 1, Exercise 1.
2. Open the properties of PriorityPrinter.
3. Click the Advanced tab.
4. Set the Priority to 99 (highest).

Exercise 3: Examine Active Directory Printer Objects

1. Open Active Directory Users And Computers.
2. From the View menu, select Users, Groups, And Computers As Containers.
3. Expand the Domain Controllers OU. Note that Server01 appears as a subcontainer.
4. Select Server01 in the tree.

   The printer objects appear in the details pane. If objects do not appear for the printers you created in Exercises 1 and 2, wait a few minutes. The print server may take a moment to publish its printers to Active Directory. You may need to press F5 (refresh) to see the printer objects once they are published.
5. Open the properties of the PriorityPrinter object.

   Note the differences between the properties that are published to Active Directory and the properties that you would see for the printer in the Printers And Faxes folder. Active Directory maintains a more limited number of properties—the properties that are most likely to be used in a search for a printer. Note also that changing a property in Active Directory does not change the property of the printer; but changing a property of the printer will, eventually, update the corresponding property in the Active Directory printer object.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're administering a Windows Server 2003 computer configured as a print server. Users in the Marketing group complain that they cannot print documents using a printer on the server. You view the permissions in the printer's properties. The Marketing group is allowed Manage Documents permission. Why can't the users print to the printer?
   a. The Everyone group must be granted the Manage Documents permission.
   b. The Administrators group must be granted the Manage Printers permission.
   c. The Marketing group must be granted the Print permission.
   d. The Marketing group must be granted the Manage Printers permission.
2. You're setting up a printer pool on a Windows Server 2003 computer. The printer pool contains three print devices, all identical. You open the properties for the printer and select the Enable Printer Pooling option on the Ports tab. What must you do next?
   a. Configure the LPT1 port to support three printers.
   b. Select or create the ports mapped to the three printers.
   c. On the Device Settings tab, configure the installable options to support two additional print devices.
   d. On the Advanced tab, configure the priority for each print device so that printing is distributed among the three print devices.
3. You're the administrator of the Windows Server 2003 computer that is configured as a print server, and you want to administer the print services from a Web browser on a client computer. The server is named Mktg1, but you don't know the share name of the printer. Which URL should you use to connect to the printer?
   a. *http://mktg1/printers*
   b. *http://printers/mktg1*
   c. *http://windows/web/printers*
   d. *http://windows/mktg1*

4. You want to configure a logical printer so that large, low-priority documents will be printed overnight. Which of the following options will you configure in the printer's Properties dialog box?

   a. Priority

   b. Available From / To

   c. Start Printing After Last Page Is Spooled

   d. Print Directly To The Printer

   e. Keep Printed Documents

Lesson Summary

The Windows printer model supports the creative and flexible utilization of printers through logical printers. You can add one logical printer that sends jobs to multiple devices (a printer pool) or multiple logical printers that send jobs to one device, with each logical printer preconfigured with printer settings, print defaults, and permissions to support a particular type of printing task.

Printers are published to Active Directory, making it easy for users to find and connect to printers. Windows Server 2003 supports printer location tracking, which further simplifies printer searches. It is even possible to administer and print to printers over the intranet or Internet using IPP.

Lesson 3: Maintaining, Monitoring, and Troubleshooting Printers

Once logical printers have been set up, configured and shared on print servers, and once clients have been connected to those printers, you must begin to maintain and monitor those logical and physical printers. This lesson will give you guidance in the maintenance and troubleshooting of printers in a Windows Server 2003 environment. You will learn to support printer drivers, to redirect printers, to configure performance and utilization logs, and to methodically troubleshoot print errors.

After this lesson, you will be able to

- Manage printer drivers
- Redirect a printer
- Monitor printer performance
- Audit printer access
- Troubleshoot printer failures

Estimated lesson time: 20 minutes

Maintaining Printers

There are no regular maintenance tasks for the print service on a Windows Server 2003 computer. The maintenance tasks defined below are typically performed on a periodic, as-needed basis. Keep in mind that when managing printers, actions may affect an entire printer or all printers on the print server, not just individual print jobs.

Managing Printer Drivers

The first grouping of maintenance tasks relate to drivers on the print server. As mentioned earlier in the lesson, it is helpful to install drivers for all client platforms that will use a particular shared printer. Windows clients will download the driver automatically when they connect to the printer. Drivers for various platforms are installed by clicking Additional Drivers on the Sharing tab of a printer's Properties dialog box.

To update drivers for a single logical printer, select the Advanced tab of the Properties dialog box and click New Driver. You will then be able to select additional drivers by indicating the manufacturer and model, or by clicking Have Disk and providing the manufacturer's drivers.

You can also manage drivers for the print server as a whole. In the Printers And Faxes folder, select Server Properties from the File menu and click the Drivers tab. Here you can add, remove, reinstall, or access the properties of each of the drivers on the print server. Changes made to these drivers will affect all printers on the server.

If you want to list all of the files related to a particular printer driver, open the print server's Drivers tab select the driver, and click Properties. The names and descriptions of all the files that are part of the specific driver will appear. From this list, it is possible to view details regarding any of the files by selecting the file and then clicking Properties.

Redirecting Print Jobs

If a printer is malfunctioning, you can send documents in the queue for that printer to another printer connected to a local port on the computer, or attached to the network. This is called *redirecting* print jobs. It allows users to continue sending jobs to the logical printer, and prevents users with documents in the queue from having to resubmit the jobs.

To redirect a printer, open the printer's Properties dialog and click the Ports tab. Select an existing port or add a port. The check box of the port of the malfunctioning printer is immediately cleared unless printer pooling is enabled, in which case you must manually clear the check box.

Because print jobs have already been prepared for the former printer, the printer on the new port must be compatible with the driver used in the logical printer. All print jobs are now redirected to the new port. You cannot redirect individual documents. In addition, any documents currently printing cannot be redirected.

Monitoring Printers

Windows Server 2003 provides several methods to monitor printers and printing resources.

Using System Monitor and Performance Logs and Alerts

The System Monitor and Performance Logs And Alerts snap-ins, both of which are included in the Performance MMC, allow you to observe real-time performance of printers, log metrics for later analysis, or set alert levels and actions. System Monitor and Performance Logs And Alerts are discussed in detail in Chapter 12. To add a counter to System Monitor, right-click the graph area and choose Add Counters. Select the performance object (in this case Print Queue), the desired counters, and the instance representing the logical printer to monitor.

After selecting Print Queue as the performance object, a list of all available performance counters is provided. You can select any counter and click Explain to learn about that particular performance metric.

The most important performance counters for monitoring printing performance are the following:

- **Bytes Printed/Sec** The number of bytes of raw data per second that are sent to the printer. Low values for this counter can indicate that a printer is underutilized, either because there are no jobs, print queues are not evenly loaded, or the server is too busy. This value varies according to the type of printer. Consult printer documentation for acceptable printer throughput values.
- **Job Errors** Number of job errors. Job errors are typically caused by improper port configuration; check port configuration for invalid settings. A printing job instance will increment this counter only once, even if it happens multiple times. Also, some print monitors do not support job error counters, in which case the counter will remain at 0.
- **Jobs** The number of jobs being spooled.
- **Total Jobs Printed** The number of jobs sent to the printer since the spooler was started.
- **Total Pages Printed** The number of pages printed since the spooler was started. This counter provides a close approximation of printer volume, although it may not be perfect, depending on the type of jobs and the document properties for those jobs.

Exam Tip The Total Jobs Printed and Total Pages Printed counters are cumulative. They represent the number of jobs or pages printed since the system was started or since the spooler was restarted.

Using System Log

Using Event Viewer, you can examine the System log as a source of information regarding spooler and printer activity. By default, the spooler registers events regarding printer creation, deletion, and modification. You will also find events containing information about printer traffic, hard disk space, spooler errors, and other maintenance issues.

To control or modify spooler event logging, open the Printers And Faxes folder and choose Server Properties from the File menu. Click the Advanced tab to access the properties as shown in Figure 8-9. From this page, you can control printer event log entries and print job notifications. This is also the tab that enables you to move the print spooler folder—an important task when configuring an active print server, or when an existing print spool folder's disk volume becomes full.

Figure 8-9 The Advanced tab of the Print Server Properties dialog box

Auditing Printer Access

Printer access, like file and folder access, can be audited. You can specify which groups or users and which actions to audit for a particular printer. After enabling object access auditing policy, you can view resulting audit entries using Event Viewer.

To configure auditing for a printer, open its Properties dialog box, click the Security tab, and then click Advanced. Click the Auditing tab and add entries for specific groups or users. For each security principal you add to the audit entry list, you can configure auditing for successful or failed access based on the standard printer permissions, including Print, Manage Documents, and Manage Printers.

You must then enable the Audit Object Access policy, which is located in group or local policy under Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy. After the policy has taken effect, you can examine the Security event log to see and analyze entries made based on printer auditing.

Tip Printer auditing creates dozens of entries for a single print job. It is therefore only useful when troubleshooting very specific problems. Printer auditing should not be used to monitor use or to bill for printer usage. Instead, performance counters such as Total Jobs Printed or Total Pages Printed should be analyzed.

Troubleshooting Printers

Troubleshooting is an important part of printer management. The following guidance will help you understand, identify, and address the types of incidents and problems that may occur in Windows Server 2003 printing.

Remember when troubleshooting that printing includes multiple components, typically:

- The application that is attempting to print.
- The logical printer on the computer on which the application is running.
- The network connection between the print client and the shared logical printer on the server.
- The logical printer on the server—its spool, drivers, security settings, and so on.
- The network connection between the print server and the printer.
- The printer itself—its hardware, configuration, and status.

An efficient way to solve most problems associated with printing is to troubleshoot each component logically and methodically.

Identify the Scope of Failure

If the user can print a job from another application on his or her computer, the error is most likely related to the failed job's application, rather than with the computer, the network, the print server, or the printer hardware. However, in some cases, using a different driver or data type can solve an application's print errors.

If the user cannot print to the printer from any application, identify whether the user can print to other printers on the same print server, or on other print servers. If all possibilities fail, and if other users can print to the printers on the network, the error is likely localized to the user's computer.

Try creating a local printer on the problematic system that points directly to the printer's port. In other words, bypass the printer server. If this process succeeds, there is a problem on the print server, with communication between the user's system and the print server, or with the printer connections on the client.

Verify That the Print Client Can Connect to the Print Server

You can confirm connectivity between the print client and the print server by opening the printer window from the Printers And Faxes folder on the client computer. If the printer window opens, showing any documents in the printer queue, the client is successfully connecting to the shared printer. An error opening the printer window would indicate a potential networking, authentication, or security permissions problem. Attempt to ping the print server's IP address. Click Start, choose Run, and type ***printserver***.

If the window opens showing the Printers And Faxes folder and any shared folders, the client is connecting to the server. Double-check security permissions on the logical printer.

Verify That the Printer Is Operational

Check the printer itself and ensure that it is in the ready state (ready to print). Print a test page from the printer console. Check the cable connecting the printer to the print server or the network. If the printer is network attached, confirm that the network interface card light is on, indicating network connectivity.

Verify That the Printer Can Be Accessed from the Print Server

Most printers can display their IP address on the printer console or by printing out a configuration page. Confirm that the printer's IP address matches the IP address of the logical printer's port. The port's IP address can be seen in the printer's Properties dialog box on the Ports tab. Ensure that it is possible to communicate with the printer over the network by pinging the printer's IP address.

Verify That the Print Server's Services Are Running

Using the Services MMC, check that services required for the printer are working properly. For example, confirm that the remote procedure call (RPC) service is running on the print server. RPC is required for standard network connections to shared printers. Confirm also that the print spooler service is running on the print server.

Tip The Net Stop Spooler command and Net Start Spooler command can be executed from the command prompt to restart the print spooler service. If you restart the spooler using command-line or user interface methods, all documents in all printer queues on the server are deleted.

You can also examine the volume on which the spool folder is stored to ensure that there is sufficient disk space for spooling. The spool folder location can be discovered and modified in the Server Properties dialog box, which you can access by choosing Server Properties from the File menu of the Printers And Faxes folder.

Note By default, the spool folder points to %*Systemroot*%\System32\Spool\Printers. For a high-volume print server, consider moving the spool folder to a partition other than the system or boot partition. If the partition where the spool folder resides fills to capacity with print jobs, printing will stop and, more importantly, the operating system might become unstable.

You should also look at the System log to see if the spooler has registered any error events, and, in the Printers And Faxes Folder, make sure that the printer is not in Offline mode.

Attempt to print a job from an application on the print server. If you can print to the printer from the print server, the problem is not with the printer. If you cannot print to the printer from an application on the print server, create a new printer directed at the same port and attempt to print to the new printer. If that job succeeds, there is a problem in the configuration of the original logical printer. If that job is unsuccessful, there is a problem communicating with the printer, or with the hardware itself.

Practice: Troubleshooting a Printer

In this practice, you will redirect a printer. Redirecting a printer is useful in both proactive and reactive troubleshooting. If you are going to take a printer offline, you can redirect its logical printer(s) to another device that is compatible with the logical printer's driver. If a printer fails due to a paper jam or other error, you can also redirect the jobs that have already been sent to, and spooled by, the logical printer, so that users do not have to wait for the failed printer to be repaired, and do not have to resubmit their jobs.

Note that additional troubleshooting practice is included in the "Case Scenario Exercise" and "Troubleshooting Lab" sections of this chapter.

Exercise 1: Redirect a Printer

If a printing device fails, you can redirect print jobs to another printer. Assume you are printing to HPLJ8100. While your job is in the queue, a job ahead of yours encounters a paper jam.

1. Open the Printers And Faxes folder and ensure that HPLJ8100 is offline. If it is not, right-click the printer and choose Use Printer Offline. This will prevent generating errors because the printer is directed to a non-existent network port.
2. Open Notepad and enter text into the blank document.
3. Choose the Print command from the File menu and select HPLJ8100 as the printer.
4. In the Printers And Faxes folder, double-click HPLJ8100 to open its printer window. Confirm that your print job is in the queue.
5. From the Printer menu, choose Properties.
6. Click the Ports tab.
7. As it was configured in Lesson 1, the printer should use the network port IP_10.0.0.51.

8. Select the check box next to the port IP_10.0.0.52.
9. Click OK. You have now redirected the printer. All jobs in the queue, except any in-progress jobs, will be directed to the new port. The printer attached to the new port must be compatible with the driver used by this logical printer, because jobs have already been processed and spooled based on the existing driver.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. A Windows 2003 Server is configured as a print server. In the middle of the workday, the printer fuse fails, and must be replaced. Users have already submitted jobs to the printer, which uses IP address 192.168.1.81. An identical printer uses address 192.168.1.217, and is supported by other logical printers on the server. What actions should you take so that users' jobs can be printed without resubmission?
   a. In the failed printer's Properties dialog box, select Enable Printer Pooling.
   b. At the command prompt, type **Net Stop Spooler**.
   c. At the command prompt, type **Net Start Spooler**.
   d. In the failed printer's Properties dialog box, select the port 192.168.1.217.
   e. In the failed printer's Properties dialog box, click Add Port.
   f. In the Printers And Faxes folder, right-click the failed printer and choose Use Offline.
2. You're setting up printing on a Windows Server 2003 computer. You attach a printer, configure a logical printer, and submit documents for printing, but the documents do not print completely and sometimes come out garbled. What is the most likely cause of the problem?
   a. There's insufficient hard disk space for spooling.
   b. You're using an incorrect printer driver.
   c. The selected port is not correct.
   d. The device settings for the printer are using an incorrect font substitution.

3. Which of the following options will give you the clearest picture of printer utilization—allowing you to understand the consumption of printer toner and paper?

   a. Configure auditing for a logical printer and audit for successful use of the Print permission by the Everyone system group.

   b. Export the System log to a comma-delimited text file and use Excel to analyze spooler events.

   c. Configure a performance log and monitor the Total Pages Printed counter for each logical printer.

   d. Configure a performance log and monitor the Jobs counter for each logical counter.

Lesson Summary

- The drivers for a logical printer can be updated or added using the properties of that printer. Drivers can be added, removed, or reinstalled for all printers on a print server using the Drivers tab of the Server Properties dialog box.
- If a printer is to be taken offline, or has already failed, you can redirect all its jobs, except those in progress, to another printer by adding or selecting the new printer's port in the properties of the original logical printer. The alternate port must represent a printer which is compatible with the driver in use by the original printer.
- The Total Jobs Printed and Total Pages Printed performance counters can help you monitor printer utilization. Bytes Printed/Sec and Errors counters will help you monitor potential problems with a printer.
- System events, logged by the spooler service, and security events, logged by enabling auditing on a printer and the Audit Object Access policy, can provide additional insight into printer functionality.
- Because the Windows Server 2003 printer model is modular, with the printer itself, the logical printer on a print server, and the printer on a client connected to the server's shared printer, you can methodically troubleshoot a printer failure by addressing each component and the links between those components.

Case Scenario Exercise

Printer usage is going through the roof at Contoso Ltd., and the chief operating officer has asked you to begin billing for printer usage by the Marketing and Sales departments, each of which are heavy users of printers.

Think Through Your Solution

1. What is the most effective way to monitor printer usage when you are billing for printer use?

   Windows Server 2003 adds a Printer Queue performance object, which allows you to monitor printer usage for each logical printer defined on the server. The Total Pages Printed counter provides important information about printer use. It is not perfect, because certain document properties and special printing features (such as a booklet printing or multiple-pages-per-page setting) will affect the printer hardware directly without the spool's being able to track their effects. However, it is the best approximation available. By configuring a performance log and capturing the counter, you can later analyze the log and bill for usage.

2. How can you monitor the Total Pages Printed counter for the Sales and the Marketing group separately?

   The Total Pages Printed counter captures performance data for a single, logical printer. To monitor the two groups separately, you must configure two separate logical printers. Each printer will address the same port—the same physical printer—but will allow only users from one group to print.

Set Up the Printers

If you are unsure how to install a logical printer, refer to Lesson 1, Exercise 1. Create two printers using the Add Printer Wizard. Use the settings described in the following tables to complete the Add Printer Wizard and the Add Standard TCP/IP Port Wizard.

Table 8-1 Sales Printer

Description	Setting
Local or Network Printer	Local printer attached to this computer. Do *not* use Plug and Play to detect the printer.
Select a Printer Port	Create a New Port: "Standard TCP/IP Port"
Printer Name or IP Address	10.0.0.53
Port Name	IP_10.0.0.53
Device Type	Hewlett Packard Jet Direct
Manufacturer	HP
Printer model	HP LaserJet 8100 Series PCL

Table 8-1 Sales Printer (Continued)

Description	Setting
Driver to use	Keep existing driver
Printer name	SalesPrinter
Default printer option	No
Share name	SalesPrinter
Location	NYC/US/1802Americas/42/B
Comment	Black and White Output Laser Printer–High Volume
Print a test page	No

Table 8-2 Marketing Printer

Description	Setting
Local or Network Printer	Local printer attached to this computer. Do *not* use Plug and Play to detect the printer.
Select a Printer Port	Use the following port: IP_10.0.0.53
Manufacturer	HP
Printer model	HP LaserJet 8100 Series PCL
Driver to use	Keep existing driver
Printer name	MarketingPrinter
Default printer option	No
Share name	MarketingPrinter
Location	NYC/US/1802Americas/42/B
Comment	Black and White Output Laser Printer–High Volume
Print a test page	No

Create Printer Users Groups

To assign permissions to the printers, you will need security groups. (If you are unsure how to create groups, refer to Chapter 4.) Create two security groups of Domain Local scope: Marketing Printer Users and Sales Printer Users.

Assign Permissions to the Printers

1. From the Printers And Faxes folder, open the Properties page of the SalesPrinter.
2. Click the Security tab.
3. Select the Everyone group and click Remove.
4. Click Add.
5. Type Sales Printer Users and click OK.

6. Assign Allow Print permission to the Sales Printer Users.
7. Repeat steps 1 through 6 to allow only the Marketing Printer Users group Print permission to the MarketingPrinter.

Configure a Performance Log

1. Open the Performance MMC from the Administrative Tools group.
2. Expand the Performance Logs And Alerts node and select Counter Logs.
3. Right-click Counter Logs and choose New Log Settings.
4. Enter the log name **Printer Utilization**.
5. Click OK. The Printer Utilization log's Properties dialog box appears.
6. Click Add Counters.
7. From the Performance Object drop-down list, select Print Queue.
8. In the Counters list, select Total Pages Printed.
9. In the Instances list, select SalesPrinter.
10. Click Add.
11. In the Instances list, select MarketingPrinter.
12. Click Add.
13. Click Close. The Printer Utilization dialog box indicates that the log will now track Total Pages Printed for each print queue.
14. Select 30 minutes as the sampling interval by typing **30** in the Interval box and selecting Minutes from the Units drop-down list.

Note Because Total Pages Printed is cumulative, from the time a print server starts, or from the time the spooler service is restarted, it is unnecessary to maintain a short sampling interval. You could sample at very long intervals as long as the server or the spooler service is not restarted in the middle of those intervals.

15. Click OK to close the Printer Utilization dialog box.
16. If you have not configured another performance log on this computer, you will be prompted to create the "C:\Perflogs" folder, in which logs are saved by default. Click Yes to confirm.
17. In the Performance Logs detail pane, the Printer Utilization log is green, indicating that it is running.
18. Stop the log by right-clicking it and choosing Stop.

Once a performance log has been created, you can examine the log in System Monitor. Click the View Log Data button on the System Monitor toolbar and you can add the performance log you generated. This particular log will not be valid for two reasons. First, two samples must be saved in a performance log for System Monitor to make use of the log's data. Unless you wait 60 minutes, or decrease the sampling interval, you will not be able to load the log. Second, Total Pages Printed will not increment because the printer does not exist, so pages do not print.

Troubleshooting Lab

The Marketing department is complaining about print quality on the MarketingPrinter. When they print from their Windows XP desktops using Microsoft Office applications, documents print perfectly. But when they print from Adobe applications, the documents do not always reflect the desired results. The Sales department, which uses a mix of Windows 2000 and Windows XP workstations, Microsoft Office, and Microsoft Customer Relationship Management (CRM), does not report any problems with the SalesPrinter.

As you consider the problem, it occurs to you that some applications produce different results depending on whether the printer is using PostScript or a non-PostScript driver.

Analyze the Solution

Where should you consider adding PostScript drivers? (Choose all that apply.)

a. The Server Properties dialog box of the print server.

b. The printer Properties dialog box of the MarketingPrinter.

c. The printer Properties dialog box of the SalesPrinter.

d. The printers installed on the desktops of each marketing department user.

The correct answer is b. Adding the PostScript driver for the MarketingPrinter will cause that printer to use the PostScript driver, without affecting the SalesPrinter. Although each client printer will require the PostScript driver as well, you do not need to add the driver manually. Windows 2000 and Windows XP clients will download the new driver automatically.

Change the Printer Driver

1. Open the Printers And Faxes folder.
2. Open the Properties dialog box of the MarketingPrinter.
3. Click the Advanced tab.
4. Click New Driver. The Add Printer Driver Wizard appears.
5. Click Next.

6. Select the Manufacturer: HP.
7. Select the Printer: HP LaserJet 8100 Series PS.
8. Click Next, and then click Finish.
9. Notice that the PostScript driver is now the default driver.
10. Click the Driver drop-down list and you will find that the former, PCL driver is still listed. If changing the driver to PostScript does not solve the problem, you can easily switch back to the PCL driver.

Chapter Summary

- Printer implementation in Windows Server 2003 is modular, consisting of the printer hardware itself, a print server with a shared, logical printer representing the physical printer by indicating that printer's local or network attached port, and a logical printer on a client that connects to the shared printer on the print server. Understanding the structure and the terminology is critical because documentation and the user interface is inconsistent and sometimes misleading.
- Shared printers are published to Active Directory, which enables users to easily search for printers based on location or other printer properties.
- When a user finds a printer in the Find Printers dialog box, double-clicking the printer installs the printer to the user's computer. Computers running the Windows operating system download the driver from the server automatically if an administrator has loaded all appropriate drivers in the shared printer.
- A single logical printer can direct jobs to more than one port, creating a printer pool.
- A single physical printer (port) can be served by multiple logical printers, each of which can configure unique properties, drivers, settings, permissions, or monitoring characteristics. Such a structure enables you to leverage printer hardware with incredible flexibility.
- Printers can be managed, installed, and printed to via the Web if Internet Printing has been installed and enabled on the print server.
- Event logs and performance counters allow you to monitor printers for potential signals of trouble, and for utilization statistics.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- The important distinction between a printer—the hardware, also known as the print device or physical printer—and a logical printer—also known as a printer.
- The difference between a printer in the Printers And Faxes folder and an Active Directory printer object.
- How to manage printer ports. Understand the difference between, and how to configure, printer pooling and printer redirection.
- How to configure multiple logical printers to a single physical printer. Be familiar with the variety of properties that can be configured uniquely in each logical printer, including security permissions.
- How to monitor printer utilization and troubleshoot printer problems.

Key Terms

Logical printer Represents a physical printer by serving the printer's port. The logical printer includes the queue, the drivers, settings, permissions, and printing defaults that manage the creation of a print job for a printer.

Network printer In the context of the Microsoft Windows user interface, a logical printer that is a client of—that is connected to—a shared logical printer on another computer. Not to be confused with a network attached printer, which is served by a *local printer* on the print server.

9 Maintaining the Operating System

Exam Objectives in this Chapter:

- Manage software update infrastructure
- Manage software site licensing

Why This Chapter Matters

In 2002, the Code Red worm and its derivatives, Code Red v2 and Code Red II, tore through the Internet, exploiting a hole in Microsoft Index Server. Although the worms themselves did not cause tremendous damage, their astounding infection rate was a wake-up call to the tens of thousands of IT professionals who had spent hours upon hours securing and updating their systems. The wake-up call was particularly poignant because Microsoft had patched the Index Server vulnerability a month before the worms wreaked their havoc. It was clearer than ever that servers and workstations must be kept current with code updates. Nor was it a wise strategy to wait for Service Pack 3 before deploying Service Pack 2, as many enterprises had done in the past. Software updates now became part and parcel of the security strategies of an organization.

In this chapter, you will learn how to apply Microsoft Software Update Services (SUS) to keep servers and desktops up to date. SUS allows an enterprise to centralize the downloading, testing, approval, and distribution of Windows-critical updates and Windows security rollups. This service will play a significant role in maintaining the integrity of your enterprise network. You will also learn how to deploy Service Packs to one or more machines. Finally, you will examine the components of site software licensing.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to administering Windows Software Update Services, service pack deployment, and licensing. Although it is advantageous to have two computers (a Microsoft Windows Server 2003 computer and a client running Windows XP or Windows 2000 Professional), you can complete the exercises in this chapter with only one computer. Prepare the following:

- A Windows Server 2003 (Standard Edition or Enterprise Edition) installed as Server01 and configured as a domain controller in the domain *contoso.com*
- A first-level organizational unit (OU) named Desktops
- Networking configured to provide Internet connectivity

Lesson 1: Software Update Services

To maintain a secure computing environment, it is critical to keep systems up to date with security patches. Since 1998, Microsoft has provided Windows Update as a Web-based source of information and downloads. With Windows XP and Windows 2000 service pack 3, Microsoft added Automatic Updates, whereby a system automatically connects to Windows Update and downloads any new, applicable patches or "hot-fixes." Although the Windows Update servers and Automatic Updates client achieve the goal of keeping systems current, many administrators are uncomfortable with either computers or users deciding which patches should be installed, because a patch might interfere with the normal functioning of a business-critical application.

The latest improvements to these technologies deliver Software Update Services (SUS). SUS is a client-server application that enables a server on your intranet to act as a point of administration for updates. You can approve updates for SUS clients, which then download and install the approved updates automatically without requiring local administrator account interaction.

In this lesson you will learn to install and administer SUS on a Windows Server 2003 computer. The following lesson will guide you through issues related to client configuration.

After this lesson, you will be able to

- Install SUS on a Windows Server 2003 computer
- Configure SUS
- Install or deploy Automatic Updates for SUS clients
- Administer SUS and Automatic Updates
- Monitor, troubleshoot, back up, and restore SUS

Estimated lesson time: 30 minutes

Understanding SUS

Since 1998, Microsoft Windows operating systems have supported Windows Update, a globally distributed source of updates. Windows Update servers interact with client-side software to identify critical updates, security rollups, and enhancements that are appropriate to the client platform, and then to download approved patches.

Administrators wanted a more centralized solution that would assure more direct control over updates that are installed on their clients. Software Update Services is a response to that need. SUS includes several major components:

- **Software Update Services, running on an Internet Information Services (IIS) server** The server-side component is responsible for synchronizing information about available updates and, typically, downloading updates from the Microsoft Internet-based Windows Update servers or from other intranet servers running SUS.
- **The SUS administration Web site** All SUS administration is Web-based. After installing and configuring SUS, administration typically consists of ensuring that the SUS server is synchronizing successfully, and approving updates for distribution to network clients.
- **Automatic Updates** The Automatic Updates client is responsible for downloading updates from either Windows Update or an SUS server, and installing those updates based on a schedule or an administrator's initiation.
- **Group Policy settings** Automatic Updates clients can be configured to synchronize from an SUS server rather than the Windows Update servers by modifying the clients' registries or, more efficiently, by configuring Windows Update policies in a Group Policy Object (GPO).

Installing SUS on a Windows Server 2003 Computer

SUS has both client and server components. The server component runs on a Windows 2000 Server (Service Pack 2 or later) or a Windows Server 2003 computer. Internet Information Services (IIS) must be installed before setting up SUS and, as you learned in Chapter 6, "Files and Folders," IIS is not installed by default on Windows Server 2003. For information about how to install IIS, see Chapter 6.

SUS is not included with the Windows Server 2003 media, but it is a free download from the Microsoft SUS Web site at *http://go.microsoft.com/fwlink/?LinkID=6930*.

Note The SUS download is not available in every localized language. However, this download determines the installation and administrative interface for the server component only. Patches for *all* locales can be made available through SUS.

After downloading the latest version of SUS, double-click the file and the installation routine will start. After you agree to the license agreement, choose Custom setup and the Setup Wizard will prompt you for the following information:

- **Choose File Locations** Each Windows Update patch consists of two components: the patch file itself and metadata that specifies the platforms and languages to which the patch applies. SUS always downloads metadata, which you will use to approve updates and which clients on your intranet will retrieve from SUS. You can choose whether to download the files themselves and, if so, where to save the updates.

Tip If you elect to maintain the update files on Microsoft Windows Update servers, Auto matic Updates clients will connect to your SUS server to obtain the list of approved updates and will then connect to Microsoft Windows Update servers to download the files. You can thereby maintain control of client updating and take advantage of the globally dispersed host ing provided by Microsoft.

If you choose the Save The Updates To This Local Folder option, the Setup Wizard defaults to the drive with the most free space, and will create a folder called SUS on that drive. You can save the files to any NT file system (NTFS) partition; Microsoft recommends a minimum of 6 gigabytes (GB) of free space.

Note The SUS partition and the system partition must be formatted as NTFS.

- **Language Settings** Although the SUS administrative interface is provided in English and a few additional languages, patches are released for all supported locales. This option specifies the localized versions of Windows servers or clients that you support in your environment.
- **Handling New Versions Of Previously Approved Updates** Occasionally, an update itself is updated. You can direct SUS to approve automatically updates that are new versions of patches that you have already approved, or you can continue to approve each update manually.
- **Ready To Install** Before installation begins, the Setup Wizard will remind you of the URL clients should point to, *http://SUS_servername*. Note this path because you will use it to configure network clients.
- **Installing Microsoft Software Update Services** The Setup Wizard installs SUS.
- **Completing the Microsoft Software Update Services Setup Wizard** The final page of the Setup Wizard indicates the URL for the SUS administration site, *http://SUS_servername/SUSAdmin*. Note this path as well, because you will administer SUS from that Web location. When you click Finish, your Web browser will start and you will be taken automatically to the SUS administration page.

Software Update Services installs the following three components on the server:

- The Software Update Synchronization Service, which downloads content to the SUS server
- An IIS Web site that services update requests from Automatic Updates clients
- An SUS administration Web page, from which you can synchronize the SUS server and approve updates

IIS Lockdown

When run on a Windows 2000 server, the SUS Setup Wizard launches the IIS Lockdown Wizard to secure IIS 5.0. Windows Server 2003 is locked down by default, so IIS Lockdown is not necessary.

If you have Web applications running on an IIS server, those applications may not function properly after SUS has been installed. You can re-enable Internet Server Application Programming Interface (ISAPI) filters and open other components that are secured by IIS Lockdown. However, due to the sensitive nature of operating system updates, you should consider running SUS on a dedicated server without other IIS applications.

Configuring and Administering SUS

You will perform three administrative tasks related to SUS: configuring SUS settings, synchronizing content and approving content. These tasks are performed using the SUS Administration Web site, shown in Figure 9-1, which can be accessed by navigating to *http://SUS_servername/SUSAdmin* with Internet Explorer 5.5 or later, or by opening Microsoft Software Update Services from the Administrative Tools programs group. The administration of SUS is entirely Web-based.

Note You may need to add Server01 to the Local Intranet trusted site list to access the site. Open Internet Explorer and choose Internet Options from the Tools menu. Click the Security tab. Select Trusted Sites and click Sites. Add Server01 and Server01.contoso.com to the trusted site list.

Note You must be a local administrator on the SUS server to administer and configure Software Update Services. This is another consideration as you review dedicating the SUS server. With a dedicated SUS server, you can delegate administration of SUS without inadvertently delegating authority over other server roles or applications.

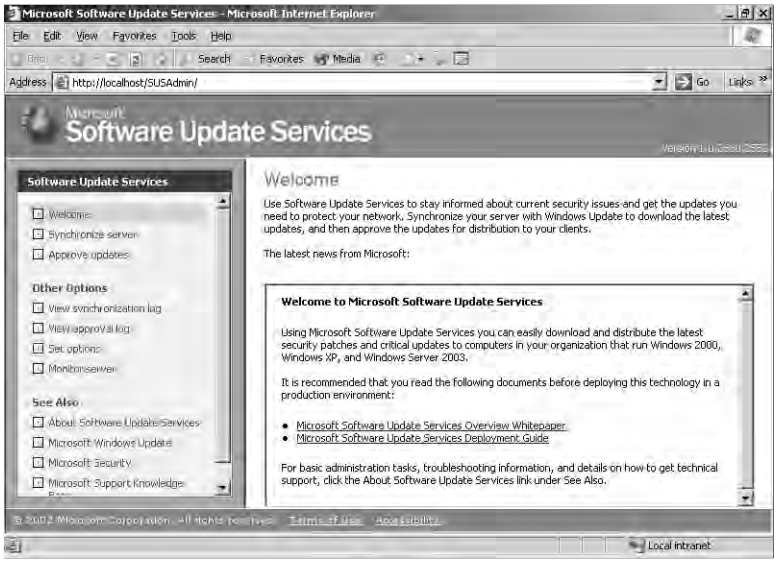

Figure 9-1 The SUS Administration Web site

Configuring Software Update Services

Although some of the configuration of SUS can be specified during a custom installation, all SUS settings are accessible from the SUS Administration Web page. From the Software Update Services administration page, click Set Options in the left navigation bar. The Set Options page is shown in Figure 9-2.

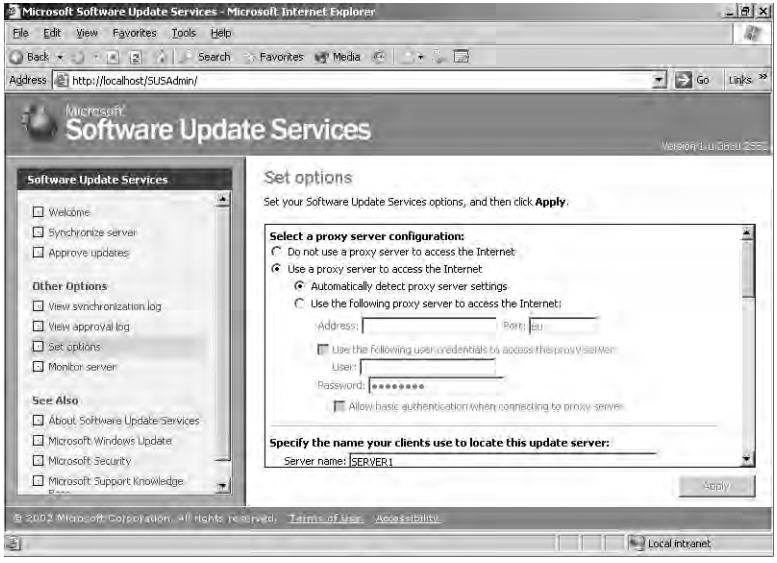

Figure 9-2 The SUS Set Options page

The configuration settings are as follows:

- **Proxy server configuration** If the server running SUS connects to Windows Update using a proxy server, you must configure proxy settings.

Tip Although the SUS server can be configured to access Windows Update through a proxy server that requires authentication, the Automatic Updates client cannot access Windows Update if the proxy server requires authentication. If your proxy server requires authentica tion, you can configure SUS to authenticate, and you must store all update content—files as well as metadata—locally.

- **DNS name of the SUS server** In the Server Name box, type the fully qualified domain name (FQDN) of the SUS server, for example, **sus1.contoso.com**.
- **Content source** The first SUS servDer you install will synchronize its content from Microsoft Windows Update. Additional SUS servers can synchronize from Windows Update, from a "parent" SUS server, or from a manually created content distribution point. See the sidebar, "SUS Topology" for more information.
- **New versions of approved updates** The Set Options page allows you to modify how SUS handles new versions of previously approved updates. This option is discussed earlier in the lesson.
- **File storage** You can modify the storage of metadata and update files. This option is also discussed earlier in the lesson.

Tip If you change the storage location from a Windows Update server to a local server folder, you should immediately perform a synchronization to download the necessary packages to the selected location.

- **Languages** This setting determines the locale specific updates that are synchronized. Select only languages for locales that you support in your environment.

Tip If you remove a locale, the packages that have been downloaded are not deleted; however, clients will no longer receive those packages. If you *add* a locale, perform a manual syn chronization to download appropriate packages for the new locale.

SUS Topology

Software Update Services is all about enabling you to control the approval and distribution of updates from Microsoft Windows Update. In a small organization, SUS can be as simple as one server, synchronizing from Windows Update and providing a list of approved updates to clients.

In a larger organization, SUS topologies can be developed to make SUS more scalable and efficient. Although the 70-290 certification exam expects you only to administer existing topologies, it is helpful to understand some of the design possibilities:

- **Multiple server topology** Each SUS server synchronizes content from Windows Update, and manages its own list of approved updates. This would be a variation of a single-server model, and each SUS server administrator would have control over that server's list of approved updates. Such a configuration would also allow an organization to maintain a variety of patch and update configurations (one per SUS server). Clients can be directed to obtain updates from an SUS server with the appropriate list of approved updates.
- **Strict parent/child topology** A "parent" SUS server synchronizes content from Windows Update and stores updates in a local folder. The SUS administrator then approves updates. Other SUS servers in the enterprise synchronize from the parent, and are configured, on the Set Options page, to Synchronize List Of Approved Items Updated From This Location (Replace Mode). This setting causes the child SUS servers to synchronize both the update files and the list of approved updates. Network clients can then be configured to retrieve updates from the SUS server in or closest to their site. In this configuration (Synchronize List Of Approved Items), administrators of child SUS servers *cannot* approve or disapprove updates; that task is managed on the parent SUS server only.
- **Loose parent/child topology** A "parent" SUS server synchronizes content from Windows Update and stores updates in a local folder. Other SUS servers in the enterprise synchronize from the parent. Unlike the strict configuration, these additional SUS servers do not synchronize the list of approved updates, so administrators of those servers can approve or disapprove updates independently. Although this topology increases administrative overhead, it is helpful when an organization wants to minimize Internet exposure (only the parent SUS server needs to connect to the Internet), and requires (as in the multiple-server model) distributed power of update approval or a variety of client patch and update configurations.

- **Test/production topology** This model allows an organization to create a testing or staging of updates. The parent SUS server downloads updates from Windows Update and an administrator approves updates to be tested. One or more clients retrieve updates from the parent SUS server and act as test platforms. Once updates have been approved, tested, and verified, the contents of the parent SUS server are copied to a manually created content distribution point on a second IIS server. Production SUS servers synchronize both the updates and the list of approved updates from the manual content distribution point. The steps for configuring such a manual distribution point are detailed in the Software Update Service Deployment White Paper, available from the Microsoft SUS Web site.

Synchronizing SUS

On the SUS Administration Web page, click Synchronize Server. On the Synchronize Server page, as shown in Figure 9-3, you can start a manual synchronization or configure automatic, scheduled synchronization. Click Synchronize Now and, when synchronization is complete, you will be informed of its success or failure, and, if the synchronization was successful, you will be taken to the Approve Updates page.

Figure 9-3 The Synchronize Server page

To schedule synchronization, click Synchronization Schedule. You can configure the time of day for synchronization, as shown in Figure 9-4, and whether synchronization occurs daily or weekly on a specified day. When a scheduled synchronization fails, SUS will try again for the Number Of Synchronization Retries To Attempt setting. Retries occur at 30-minute intervals.

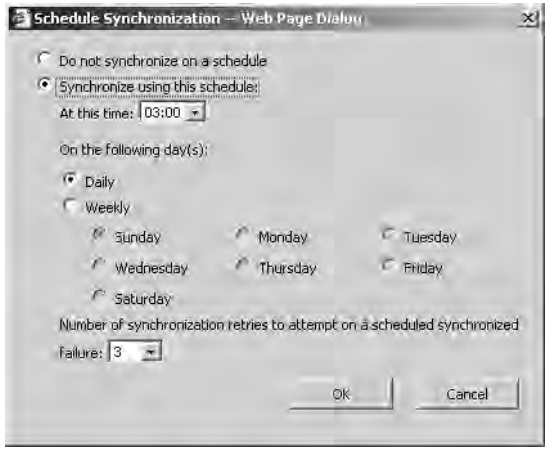

Figure 9-4 The Schedule Synchronization Web Page Dialog page

Approving Updates

To approve updates for distribution to client computers, click Approve Updates in the left navigation bar. The Approve Updates page, as shown in Figure 9-5, appears. Select the updates that you wish to approve, then click Approve. If you are unsure about the applicability of a particular update, click the Details link in the update summary. The Details page that opens will include a link to the actual *.cab file that is used to install the package, and a link to the Read More page about the update, which will open the Microsoft Knowledge Base article related to the update.

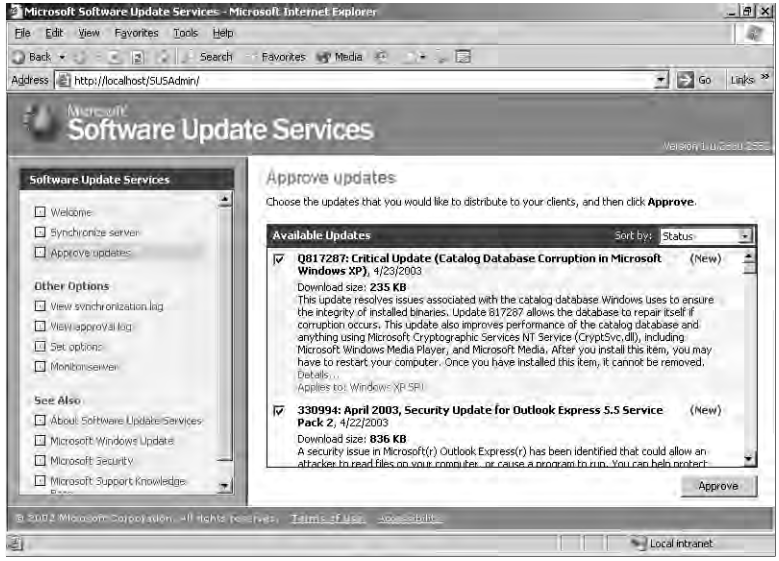

Figure 9-5 The Approve Updates page

Tip The first synchronization will download dozens of updates. It may be tedious to scroll and click each check box for approval. Instead, after clicking the first check box, press TAB twice to navigate to the next check box, and press the spacebar to select (or clear) the item.

The Automatic Updates Client

The client component of SUS is Windows Automatic Updates, which is supported on Windows 2000, Windows XP, and Windows Server 2003. The Automatic Updates client is included with Windows Server 2003, Windows 2000 Service Pack 3, and Windows XP Service Pack 1.

For clients running earlier releases of the supported platforms, you can download Automatic Updates as a stand-alone client from the Microsoft SUS Web site, at *http://go.microsoft.com/fwlink/?LinkID=6930*. The client, provided as an .msi file, can be installed on a stand-alone computer or by means of Group Policy (assign the package in the Computer Configuration\Software Settings policy), SMS, or even a logon script. If a localized version of the client is not available, install the English version on any locale.

The Automatic Updates client of Windows Server 2003 is configured to connect automatically to the Microsoft Windows Update server and download updates, then prompt the user to install them. This behavior can be modified by accessing the Automatic Updates tab in the System Properties dialog box, accessible by clicking System in Control Panel, in Windows XP and Windows Server2003. In Windows 2000 click Automatic Updates in Control Panel. The Automatic Updates tab is shown in Figure 9-6. Automatic Updates can also be configured using GPOs or registry values.

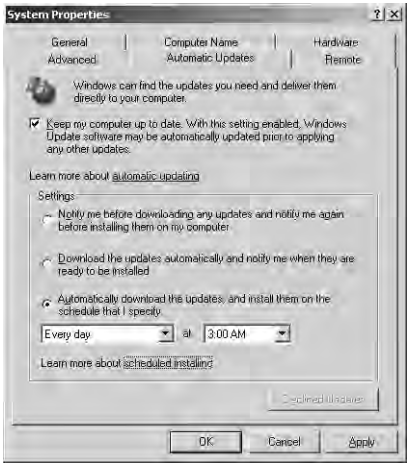

Figure 9-6 The Automatic Updates tab of the System Properties dialog box

Download Behavior

Automatic Updates supports two download behaviors:

- **Automatic** Updates are downloaded without notification to the user.
- **Notification** If Automatic Updates is configured to notify the user before downloading updates, it registers the notification of an available update in the system event log and to a logged-on administrator of the computer. If an administrator is not logged on, Automatic Updates waits for a user with administrator credentials before offering notification by means of a balloon in the notification area of the system tray.

Once update downloading has begun, Automatic Updates uses the Background Intelligent Transfer Service (BITS) to perform the file transfer using idle network bandwidth. BITS ensures that network performance is not hindered due to file transfer. All patches are checked by the SUS server to determine if they have been correctly signed by Microsoft. Similarly, the Automatic Updates client confirms the Microsoft signature and also examines the cyclical redundancy check (CRC) on each package before installing it.

Installation Behavior

Automatic Updates provides two options for installation:

- **Notification** Automatic Updates registers an event in the system log indicating that updates are ready for installation. Notification will wait until a local administrator is logged on before taking further action. When an administrative user is logged on, a balloon notification appears in the system tray. The administrator clicks the balloon or the notification icon, and then may select from available updates before clicking Install. If an update requires restarting the computer, Automatic Updates cannot detect additional updates that might be applicable until after the restart.
- **Automatic (Scheduled)** When updates have been downloaded successfully, an event is logged to the system event log. If an administrator is logged on, a notification icon appears, and the administrator can manually launch installation at any time until the scheduled installation time.

  At the scheduled installation time, an administrator who is logged on will be notified with a countdown message prior to installation, and will have the option to cancel installation, in which case the installation is delayed until the next scheduled time. If a non-administrator is logged on, a warning dialog appears, but the user cannot delay installation. If no user is logged on, installation occurs automatically. If an update requires restart, a five-minute countdown notification appears informing users of the impending restart. Only an administrative user can cancel the restart.

Tip If a computer is not turned on at the scheduled Automatic Updates installation time, installation will wait to the next scheduled time. If the computer is never on at the scheduled time, installation will not occur. Ensure that systems remain turned on to be certain that Auto matic Updates install successfully.

Configuring Automatic Updates Through Group Policy

The Automatic Updates client will, by default, connect to the Microsoft Windows Update server. Once you have installed SUS in your organization, you can direct Automatic Updates to connect to specific intranet servers by configuring the registry of clients manually or by using Windows Update group policies.

To configure Automatic Updates using GPOs, open a GPO and navigate to the Computer Configuration\Administrative Templates\Windows Components\Windows Update node. The Windows Update policies are shown in Figure 9-7.

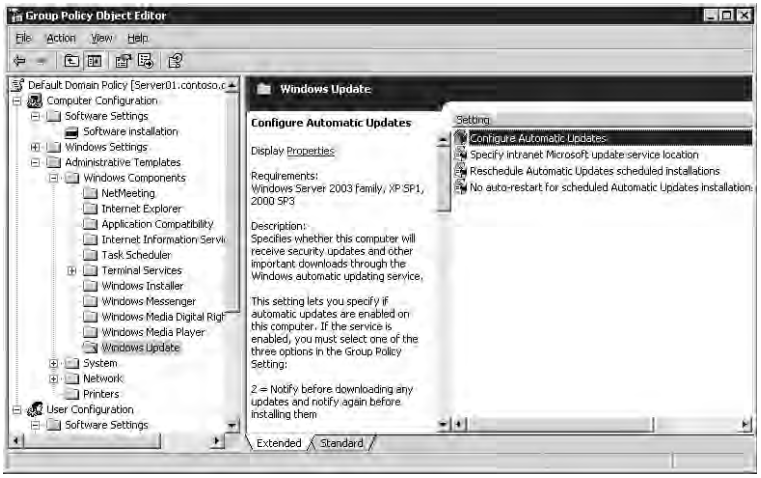

Figure 9-7 Windows Update policies

Note If you edit policy on a Windows 2000 Active Directory server, the policies may not appear. Automatic Updates policies are described by the *%Windir%*\Inf\Wuau.inf administra tive template, which is installed by default when Automatic Updates is installed. If Automatic Updates has not been installed on the domain controller to which you are connected (typi cally, the PDC Emulator), you must right-click the Administrative Templates node and choose Add/Remove Templates, click Add, then locate the Wuau.inf template, perhaps by copying it from a system that does have Automatic Updates installed.

The following policies are available, each playing an important role in configuring effective update distribution in your enterprise:

- **Configure Automatic Updates** The Configure Automatic Updates Behavior determines the behavior of the Automatic Updates client. There are three options: Notify For Download And Notify For Install, Auto Download And Notify For Install, and Auto Download And Schedule The Install. These options are combinations of the installation and download behaviors discussed earlier in the lesson.
- **Reschedule Automatic Updates Scheduled Installations** If installations are scheduled, and the client computer is turned off at the scheduled time, the default behavior is to wait for the next scheduled time. The Reschedule Automatic Updates Scheduled Installations policy, if set to a value between 1 and 60, causes Automatic Updates to reschedule installation for the specified number of minutes after system startup.
- **No Auto-Restart For Scheduled Automatic Updates Installations** This policy causes Automatic Updates to forego a restart required by an installed update when a user is logged on to the system. Instead, the user is notified that a restart is required for installation to complete, and can restart the computer at his or her discretion. Remember that Automatic Updates cannot detect new updates until restart has occurred.
- **Specify Intranet Microsoft Update Service Location** This policy allows you to redirect Automatic Updates to a server running SUS. By default, the client will log its interactions on the SUS server to which it connects. However, this policy allows you to point clients to another server running IIS for statistics logging. This dual policy provides the opportunity for clients to obtain updates from a local SUS server, but for all clients to log SUS statistics in a single location for easier retrieval and analysis of the log data, which is stored as part of the IIS log. IIS logs typically reside in *%Windir%*\System32\Logfiles\W3svc1.

Automatic Updates clients poll their SUS server every 22 hours, minus a random offset.

Any delay in patching should be treated as unacceptable when security vulnerabilities are being actively exploited. In such situations, install the patch manually so that systems do not have to wait to poll, download, and install patches.

After approved updates have been downloaded from the SUS server, they will be installed as configured—manually or automatically—at the scheduled time. If an approved update is later unapproved, that update is not uninstalled; but it will not be installed by additional clients. An installed update *can* be uninstalled manually, using the Add Or Remove Programs application in Control Panel.

SUS Troubleshooting

Although SUS works well, there are occasions that warrant monitoring and troubleshooting.

Monitoring SUS

The Monitor Server page of the SUS Administration Web site displays statistics that reflect the number of updates available for each platform, and the date and time of the most recent update. The information is summarized from the Windows Update metadata that is downloaded during each synchronization. Metadata information is written to disk and stored in memory to improve performance as systems request platform appropriate updates.

You can also monitor SUS and Automatic Updates using the following logs:

- **Synchronization Log** You can retrieve information about current or past synchronizations, and the specific packages that were downloaded by clicking View Synchronization Log in the left navigation bar. You can also use any text editor to open the (Extensible Markup Language) XML–based database (History-Sync.xml) directly from the SUS Web site's \AutoUpdate\Administration directory in IIS.
- **Approval Log** For information about packages that have been approved, click View Approval Log in the left navigation bar. Alternatively, you can open History-Approve.xml from the SUS Web site's \AutoUpdate\Administration directory in IIS.
- **Windows Update Log** The Automatic Updates client logs activity in the *%Windir%*\Windows Update.log file on the client's local hard disk.
- **Wutrack.bin** The client's interaction with SUS is logged to the specified statistics server's IIS logs, typically stored in the folder: *%Windir%*\System32\Logfiles \W3svc1. These logs, which are verbose and cryptic, are designed to be analyzed by programs, not by humans.

Exam Tip Although you should know what logs are available, and where they are located, you are not required for the 70-290 exam to be able to interpret cryptic messages or log entries. The SUS Deployment White Paper includes appendices with detailed information about event descriptions and log syntax.

SUS System Events

The synchronization service generates event log messages for each synchronization performed by the server, and when updates are approved. These messages can be viewed in the System log using Event Viewer. The events relate to the following scenarios:

- **Unable to connect** Automatic Updates could not connect to the update service (Windows Update or the computer's assigned SUS server).
- **Install ready—no recurring schedule** Updates listed in the event were downloaded and are pending installation. An administrator must click the notification icon and click Install.
- **Install ready—recurring schedule** Updates listed in the event are downloaded and will be installed at the date and time specified in the event.
- **Installation success** Updates listed in the event were installed successfully.
- **Installation failure** Updates listed in the event failed to install properly.
- **Restart required—no recurring schedule** An update requires a restart. If installation behavior is set for notification, restart must be performed manually. Windows cannot search for new updates until the restart has occurred.
- **Restart required—recurring schedule** When Automatic Updates is configured to automatically install updates, an event is registered if an update requires restart. Restart will occur within five minutes. Windows cannot search for new updates until after the restart has occurred.

Troubleshooting SUS

Software Update Services on a Windows Server 2003 computer may require the following troubleshooting steps:

- **Reloading the memory cache** If no new updates appear since the last time you synchronized the server, it is possible that no new updates are available. However, it is also possible that memory caches are not loading new updates properly. From the SUS administration site, click Monitor Server and then click Refresh.
- **Restarting the synchronization service** If you receive a message that the synchronization service is not running properly, or if you cannot modify settings in the Set Options page of the administration Web site, open the Microsoft Management Console (MMC) Services snap-in, right-click Software Update Services Synchronization Service and choose Restart.
- **Restarting IIS** If you cannot connect to the administration site, or if clients cannot connect to the SUS serve, restart the World Wide Web Publishing Service in the same manner.

If Automatic Updates clients do not appear to be receiving updates properly, open the registry of a client and ensure that the following values appear in HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\WindowsUpdate:

- **WUServer** Should have the URL of the SUS server, for example, *http://SUS_Servername*
- **WUStatusServer** Should have the URL of the same SUS server or another IIS server on which synchronization statistics are logged

And, in the AU subkey:

- **UseWUServer** Should be set to dword:00000001

SUS Backup and Recovery

As with any other server role or application, you must plan for recovery in the event of a server failure.

Backing Up SUS

To back up SUS, you must back up the folder that contains SUS content, the SUS Administration Web site, and the IIS metabase.

Exam Tip The process described to back up the IIS metabase is useful not only for backing up SUS, but for any other Web site or application running on Windows Server 2003 and IIS 6.0.

First, back up the metabase—an XML database containing the configuration of IIS. Using the MMC IIS snap-in, select the server to back up and, from the Action menu, select All Tasks, then Backup/Restore Configuration. Click Create Backup and enter a name for the backup. When you click OK, the metabase is backed up.

Then back up the following using Ntbackup or another backup utility:

- The default Web site, which is located unless otherwise configured in C:\Inetpub\Wwwroot.
- The SUS Administration Web site. SUSAdmin is, by default, a subfolder of C:\Inetpub\Wwwroot. In that event, it will be backed up when you back up the default Web site.
- The AutoUpdate virtual directory, also by default a subfolder of C:\Inetpub\Wwwroot.
- The SUS content location you specified in SUS setup or the SUS options. You can confirm the SUS content location in IIS manager by clicking Default Web Site and examining the path to the Content virtual root in the details pane.
- The metabase backup directory, *%Windir%*\System32\Inetsrv\Metaback, which contains the copy of the metabase made earlier.

See Also For more information about the Ntbackup utility, see Chapter 7.

This process of backing up the metabase, and then backing up the components of SUS, should be repeated regularly because updates will be added and approved with some frequency.

SUS Server Recovery

To restore a failed SUS server, perform the steps described below. If a certain step is unnecessary, you may skip it, but perform the remaining steps in sequence.

1. Disconnect the server from the network to prevent it from being infected with viruses.
2. Install Windows Server 2003, being sure to give the server the same name it had previously.
3. Install IIS with the same components it had previously.
4. Install the latest service pack and security fixes. If the server must be connected to the network to achieve this step, take all possible precautions to prevent unnecessary exposure.
5. Install SUS into the same folder it was previously installed.
6. Run Ntbackup to restore the most recent backup of SUS. This will include the SUS content folder, the Default Web Site, including the SUSAdmin and AutoUpdate virtual directories, and the IIS metabase backup.
7. Open the MMC IIS snap-in and select the server to restore. From the Action menu, select All Tasks, then Backup/Restore Configuration and select the backup that was just restored. Click Restore.
8. Confirm the success of your recovery by opening the SUS Administration Web site and clicking Set Options. Check that the previous settings are in place, and that the previously approved updates are still approved.

Note The preceding steps apply to Windows Server 2003 only. If you are recovering a Windows 2000-based SUS server, refer to SUS documentation for appropriate steps.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You are configuring a Software Update Services infrastructure. One server is synchronizing metadata and content from Windows Update. Other servers (one in each site) are synchronizing content from the parent SUS server. Which of the following steps is required to complete the SUS infrastructure?
   a. Configure Automatic Updates clients using Control Panel on each system
   b. Configure GPOs to direct clients to the SUS server in their sites
   c. Configure a manual content distribution point
   d. Approve updates using the SUS administration page
2. You are configuring SUS for a group of Web servers. You want the Web servers to update themselves nightly based on a list of approved updates on your SUS server. However, once in a while an administrator is logged on, performing late-night maintenance on a Web server, and you do not want update installation and potential restart to interfere with those tasks. What Windows Update policy configuration should you use in this scenario?
   a. Notify For Download And Notify For Install
   b. Auto Download And Notify For Install
   c. Auto Download And Schedule The Install
3. You want all network clients to download and install updates automatically during night hours, and you have configured scheduled installation behavior for Automatic Updates. However, you discover that some users are turning off their machines at night, and updates are not being applied. Which policy allows you to correct this situation without changing the installation schedule?
   a. Specify Intranet Microsoft Update Service Location
   b. No Auto-Restart For Scheduled Automatic Updates Installations
   c. Reschedule Automatic Updates Scheduled Installations
   d. Configure Automatic Update

Lesson Summary

- SUS is an intranet application that runs on IIS 6.0 (or on IIS 5.0 on a Windows 2000 Server) and is administered through a Web-based administration site: *http://SUS_Servername/SUSAdmin*.
- The SUS server synchronizes information about critical updates and security rollups and allows an administrator to configure approval centrally for each update. Typically, an enterprise configures SUS to download the actual update files as well.
- Automatic Updates, which runs on Windows 2000, Windows XP, and Windows Server 2003, is responsible for downloading and installing updates on the client.
- Group Policy can be used to configure Automatic Updates to retrieve patches from an SUS server rather than from the Windows Update servers. GPOs can also drive the download, installation and restart behavior of the client computers.

Lesson 2: Service Packs

Microsoft releases Service Packs to consolidate critical updates, security rollups, hotfixes, driver updates, and feature enhancements. As suggested at the beginning of this chapter, it is no longer feasible to wait until Service Pack 3 before installing Service Pack 2. You must stay current with Service Packs to maintain the security and integrity of your enterprise network. Software Update Services, discussed in the previous lesson, does not distribute service packs. To keep your network completely up to date with critical patches, you need to implement the skills covered in this lesson, which will allow you to deploy service packs by means of Group Policy.

After this lesson, you will be able to

- Download and extract a service pack
- Deploy a service pack with Group Policy–based software distribution

Estimated lesson time: 5 minutes

Downloading and Extracting Service Packs

When a service pack is released, Microsoft makes it available for installation and download from the Microsoft Web site. A service pack can be installed directly from a Microsoft server, in which case the client launches the service pack setup from the Microsoft site and a small setup utility is downloaded to the client. That setup utility reconnects to the Microsoft server and controls the download and installation of the entire service pack. Service packs are generally sizeable, so performing this task machine-by-machine is not an efficient deployment strategy in all but the smallest environments.

Service packs can also be obtained on CD from Microsoft and through many Microsoft resources, such as TechNet and MSDN. Service Pack CDs often include extras, such as updated administrative tools, new policy templates, and other value-added software. In an enterprise environment, it is therefore recommended to obtain the service pack media.

When you do not have access to a CD containing the service pack, and you want to deploy the service pack to more than one system, you can download the entire service pack as a single file, again from the Microsoft Web site. The service pack executable, if launched (by double-clicking, for example), triggers the installation of the service pack. This single-file version of the executable can also be *extracted* into the full folder and file structure of the service pack, just as it would be on the service pack CD, but without the value adds.

To extract a service pack, launch the executable from a command prompt with the -x switch. For example, to extract Windows XP Service Pack 1, type **xpsp1.exe -x**. You will then be prompted for a folder to which the service pack is extracted. Once the process is complete, you will see the full service pack folder structure contained in the target folder. You can then launch installation of the service pack, just as from the CD, by double-clicking I386\Update\Update.exe.

Deploying Service Packs with Group Policy

Service Pack installation requires administrative credentials on the local computer, unless the service pack is installed via Group Policy or Systems Management Server (SMS). Because service packs apply to systems, it is necessary to assign the service pack through computer-based, rather than user-based, group policy.

To distribute a service pack, create a shared folder and either extract the service pack to that folder or copy the contents of the service pack CD to the folder. Then, using the Active Directory Users And Computers snap-in, create or select an existing GPO. Click Edit and the Group Policy Object Editor console appears, focused on the selected GPO.

Expand the Computer Configuration\Software Settings node. Right-click Software Installation and choose New, then Package. Enter the path to the service pack's Update.msi file. Be certain to use a UNC format (for example, \\Server\Share) and *not* a local volume path, such as *Drive*:*Path*. In the Deploy Software dialog box, select Assigned. Close the Group Policy Object Editor console. Computers within the scope of the GPO—in the site, domain, or OU branch to which the policy is linked—automatically deploy the service pack at the next startup.

Tip Windows XP systems with Logon Optimization configured may require two restarts. Logon Optimization can be disabled by enabling the policy Always Wait For The Network At Computer Startup And Logon, found in the policy path Computer Configuration\Administrative Templates\System\Logon.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What command should you use to unpack the single file download of a service pack?
   - **a.** Setup.exe -u
   - **b.** Update.exe -x
   - **c.** Update.msi
   - **d.** <*Servicepackname*>.exe -x
2. What type of Group Policy software deployment should be used to distribute a service pack?
   - **a.** Published in the Computer Configuration Software Settings
   - **b.** Assigned in the Computer Configuration Software Settings
   - **c.** Published in the User Configuration Software Settings
   - **d.** Assigned in the User Configuration Software Settings

Lesson Summary

- Service packs can be extracted using the -x switch.
- Group Policy can deploy service packs by assigning Update.msi through the computer configuration's software settings policy.

Lesson 3: Administering Software Licenses

The End User License Agreement (EULA) is more than just a nuisance that you must click through to begin installing a new operating system, update, or application. The EULA is a binding contract that gives you the legal right to use a piece of software. In an enterprise environment, managing software licenses is critically important. In this lesson, you will learn to use the licensing tools provided by Windows Server 2003 to register and monitor licenses and compliance.

After this lesson, you will be able to

- Understand Per Server and Per Device or Per User licensing modes
- Configure licenses using the Licensing properties in Control Panel, and the Licensing administrative tool
- Create license groups

Estimated lesson time: 20 minutes

Note The Evaluation Edition of Windows Server 2003, Enterprise Edition, included on the Supplemental CD-ROM with this book, does not support licensing. You will not be able to follow along with the examples in this lesson without purchasing the full retail version of the product.

Obtaining a Client Access License

The server license for Windows Server 2003 enables you to install the operating system on a computer, but you need a Client Access License (CAL) before a user or device is legally authorized to connect to the server. CALs are obtained in bundles, and are often but not always included in the purchase of the operating system. Keep copies of the CAL certificates and your EULAs on file, in the event that your organization is audited for licensing compliance.

Tip Remember that when upgrading a server from Windows NT 4 or Windows 2000 to Windows Server 2003, you must purchase CAL upgrades as well.

You must have a CAL for any connection to a Windows Server 2003 computer that uses server components, which include file and print services or authentication. Very few server applications run so independently that the client/server connection does not require a CAL. The most significant exception to the CAL requirement is unauthenticated access conducted through the Internet. Where there is no exchange of credentials during

Internet access, such as users browsing your public Web site, no CAL is required. CALs are therefore not required for Windows Server 2003 Web Edition.

There are two types of CALs: Windows Device CALs, which allow a device to connect to a server regardless of the number of users who may use that device; and Windows User CALs, which allow a user to connect to a server from a number of devices. Windows Device CALs are advantageous for an organization with multiple users per device, such as shift workers. Windows User CALs make most sense for an organization with employees that access the network from multiple or unknown devices.

Note The licensing tools and the user interface do not yet distinguish between Windows User or Windows Device CALs. A device CAL is registered indirectly, using license groups.

The number of CALs you require, and how you track those licenses, depends on which client access licensing mode you pursue.

Per-Server Licensing

Per-server licensing requires a User or Device CAL for each concurrent connection. If a server is configured with 1,000 CALs, the 1,001st concurrent connection is denied access. CALs are designated for use on a particular server, so if the same 1,000 users require concurrent connections to a second server, you must purchase another 1,000 CALs.

Per server licensing is advantageous only in limited access scenarios, such as when a subset of your user population accesses a server product on very few servers. Per server licensing is less cost-effective in a situation where multiple users access multiple resources on multiple servers. If you are unsure which licensing mode is appropriate, select Per Server. The license agreement allows a no-cost, one-time, one-way conversion from Per Server to Per Device or Per User licensing when it becomes appropriate to do so.

Per-Device or Per-User Licensing

The Per Device or Per User licensing mode varies from the Per Seat scheme of previous versions of Windows. In this new mode, each device or user that connects to a server requires a CAL, but with that license the device or user can connect to a number of servers in the enterprise. Per User or Per Device mode is generally the mode of choice for distributed computing environments in which multiple users access multiple servers.

For example, a developer who uses a laptop and two desktops would require only one Windows User CAL. A fleet of 10 Tablet PCs that are used by 30 shift workers would require only 10 Windows Device CALs.

The total number of CALs equals the number of devices or users, or a mixture thereof, that access servers. CALs can be reassigned under certain, understandable conditions—for example a Windows User CAL can be reassigned from a permanent employee to a temporary employee while the permanent employee is on leave. A Windows Device CAL can be reassigned to a loaner device while a device is being repaired.

Per Server and Per Device or Per User licensing modes are illustrated in Table 9-1.

Table 9-1 CAL Licensing Modes

Per Server	Per User or Per Device
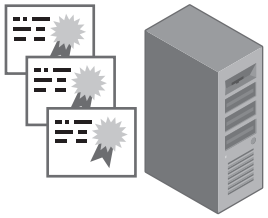	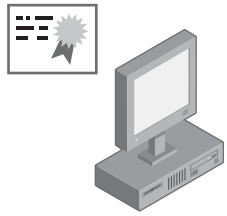
■ Traditionally licensed in Per Server mode when there are few servers that require limited access. ■ The number of CALs needed is determined by the number of concurrent connections that are required.	■ Traditionally licensed in Per User or Per Device mode when there are many servers that require frequent and widespread access. ■ Usually more economical when the number of CALs needed is determined by the number of users or devices, or both, that require access to the servers.

Tip Windows Server 2003 includes Terminal Services, also known as Remote Desktop. Remote Desktop includes a two (concurrent) connection license for administrators to connect to a remote server. For Terminal Services to perform as an application server, allowing non-administrative users to connect to hosted applications, you must acquire Terminal Services CALs, which are included in Windows XP Professional.

There are two utilities that will help you track and manage software licensing:

- **Licensing in Control Panel** The Control Panel Choose Licensing Mode tool, as shown in Figure 9-8, manages licensing requirements for a single computer running Windows Server 2003. You can use Licensing to add or remove CALs for a server running in per-server mode; to change the licensing mode from Per Server to Per Device or Per User; or to configure licensing replication.

Figure 9-8 The Choose Licensing Mode tool in Control Panel

- **Licensing in Administrative Tools** The Licensing administrative tool, discussed in the next section, allows you to manage licensing for an enterprise by centralizing the control of licensing and license replication in a site-based model.

Administering Site Licensing

The License Logging service, which runs on each Windows Server 2003 computer, assigns and tracks licenses when server resources are accessed. To ensure compliance, licensing information is replicated to a centralized licensing database on a server in the site. This server is called the site license server. A site administrator, or an administrator for the site license server, can then use the Microsoft Licensing tool in Administrative Tools program group to view and manage licensing for the entire site. This new license tracking and management capability incorporates licenses not just for file and print services, but for IIS, for Terminal Services, and for BackOffice products such as Exchange or SQL Server.

The Site License Server

The site license server is typically the first domain controller created in a site. To find out what server is the license server for a site, open Active Directory Sites And Services, expand to select the Site node then right-click Licensing Site Settings and choose Properties. The current site license server is displayed, as shown in Figure 9-9.

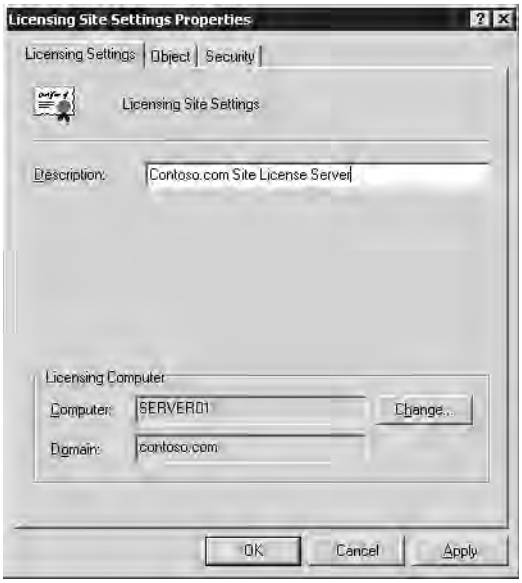

Figure 9-9 Identifying and changing the site license server

To assign the site license server role to another server or domain controller, click Change and select the desired computer. To retain the licensing history for your enterprise, you must immediately after transferring the role stop the License Logging service on the new license server, then copy the following files from the old to the new licensing server:

- %*Systemroot*%\System32\Cpl.cfg contains the purchase history for your organization.
- %*Systemroot*%\Lls\Llsuser.lls contains user information about the number of connections.
- %*Systemroot*%\Lls\Llsmap.lls contains license group information.

After all files have been copied, restart the License Logging service.

Administering Site Licenses

Once you have identified the site license server for a site, you can view the licensing information on that server opening Licensing from the Administrative Tools program group. The Server Browser tab in Licensing (as shown in Figure 9-10) enables you to manage licensing for an entire site or enterprise.

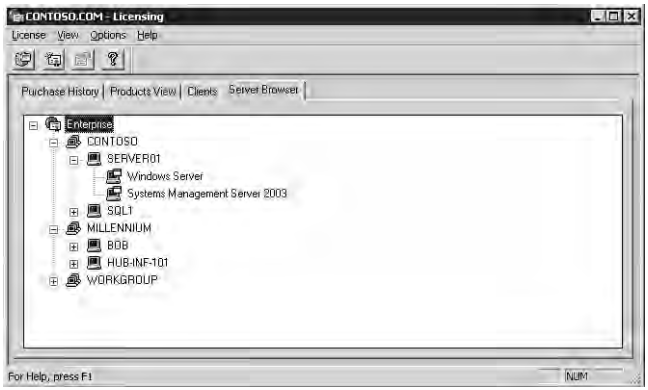

Figure 9-10 The Server Browser tab of the Microsoft Licensing administrative tool

The Server Browser page of Licensing allows you to manage any server in any site or domain for which you have administrative authority. You can locate a server and, by right-clicking it and choosing Properties, manage that server's licenses. For each server product installed on that server, you can add or remove per-server licenses. You can also, where appropriate, convert the licensing mode. Remember that per server licensing mode issues a license when a user connects to the server product. When a user disconnects from the server product, the License Logging service makes the license available to another user.

The server properties also allow you to configure license replication, which can be set on a server using its Licensing properties in Control Panel. By default, license information is replicated from a server's License Logging Service to the site license server every 24 hours, and the system automatically staggers replication to avoid burdening the site licensing server. If you want to control replication schedules or frequency, you must manually vary the Start At time and Start Every frequency of each server replicating to a particular site license server.

To manage Per Device or Per User licensing, click Licensing from the Administrative Tools program group, then choose the New License command from the License menu. In the New Client Access License dialog box, select the server product and the number of licenses purchased. Licenses are added to the pool of licenses. As devices or users connect to the product anywhere in the site, they are allocated licenses from the pool, with one license for each device or user. After a pool of licenses is depleted, license violations occur when additional devices or users access the product.

The Purchase History tab in Licensing (as shown in Figure 9-11) provides a historical overview of licenses purchased for a site, as well as the quantity, date, and administrator associated with the addition or removal of licenses.

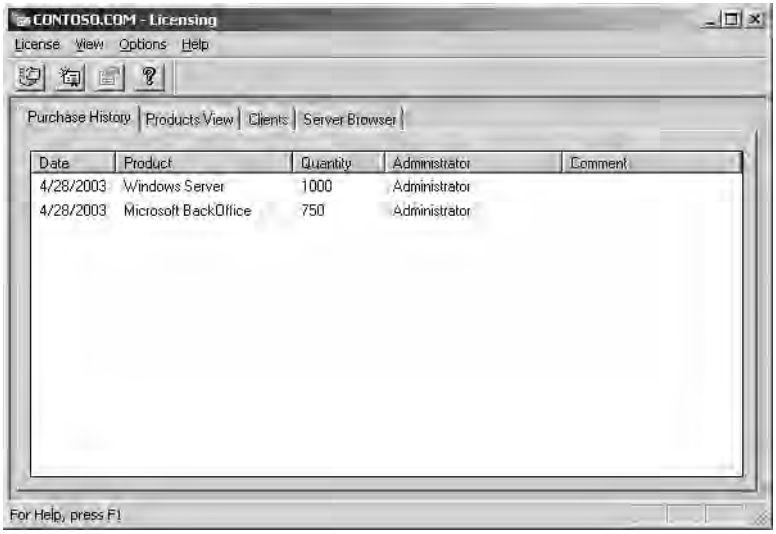

Figure 9-11 The Purchase History tab of the Microsoft Licensing administrative tool

To view cumulative information about licensing and compliance, click the Products View tab. This tab shows how many licenses have been purchased and allocated to users or devices (in Per Device or Per User mode) or the number of licenses purchased for all servers in the site and the peak connections reached to date (in Per Server mode). You can also determine compliance using the licensing status symbols shown in Table 9-2.

Table 9-2 Licensing Status Symbols

Symbol	Licensing Status
	The product is in compliance with legal licensing requirements. The number of connections is less than the number of licenses purchased.
	The product is not in compliance with legal licensing requirements. The number of connections exceeds the number of licenses purchased.
	The product has reached the legal limit. The number of connections equals the number of licenses purchased. If additional devices or users will connect to the server product, you must purchase and log new licenses.

License Groups

Per Device or Per User licensing requires one CAL for each device. However, the License Logging service assigns and tracks licenses by user name. When multiple users share one or more devices, you must create license groups, or else licenses will be consumed too rapidly.

A license group is a collection of users who collectively share one or more CALs. When a user connects to the server product, the License Logging service tracks the user by name, but assigns a CAL from the allocation assigned to the license group. The concept is easiest to understand with examples:

- **10 users share a single handheld device for taking inventory.** A license group is created with the 10 users as members. The license group is assigned one CAL, representing the single device they share.
- **100 students occasionally use a computer lab with 10 computers.** A license group is created with the 100 students as members, and is allocated 10 CALs.

To create a license group, click the Options menu and, from the Advanced menu, choose New License Group. Enter the group name and allocate one license for each client device used to access the server. The number of licenses allocated to a group should correspond to the number of devices used by members of the group.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What are the valid licensing modes in Windows Server 2003? Select all that apply.
   a. Per User
   b. Per Server
   c. Per Seat
   d. Per Device or Per User

2. You are hiring a team to tackle a software development project. There will be three shifts of programmers, and each shift will include six programmers. Each programmer uses four devices to develop and test the software, which authenticates against a Windows Server 2003 computer. What is the minimum number of CALs required if the servers involved are in Per Device or Per User licensing mode?
   a. 6
   b. 4
   c. 18
   d. 24
3. What tool will allow you to identify the site license server for your site?
   a. Active Directory Domains And Trusts
   b. The Licensing tool in Control Panel
   c. Active Directory Sites And Services
   d. DNS
4. You manage the network for a team of 500 telephone sales representatives. You have 550 licenses configured in Per Device or Per User licensing mode. A new campaign is launched and you will hire another shift of 500 reps. What do you need to do to most effectively manage license tracking and compliance?
   a. Revoke the licenses from the existing clients
   b. Delete the existing licenses, then add 500 licenses
   c. Create license groups
   d. Convert to Per Server licenses

Lesson Summary

- Windows Server 2003 provides a new mode of licensing, whereby a user can access a server product from multiple devices using one license, or a group of users can access a server product from a single device. This is called Per Device or Per User licensing.
- When more than one user accesses a server product from shared devices, add those users as a license group, and allocate licenses to that group equivalent to the number of devices.
- License information is replicated, by default every 24 hours, to the site license server.
- Licensing can be managed using the Licensing tool in Control Panel or, more centrally, using the Licensing administrative tool, from the Administrative Tools program group.

Case Scenario Exercise

You are configuring an update strategy for a network consisting of 1,000 clients running a mix of Windows XP and Windows 2000. Your goal is to prevent users from downloading updates directly from Windows Update, and to create a structure in which you can approve critical patches and security rollups for distribution.

You have recently purchased desktops and laptops, and you have applied the corporate standard image to those systems. Unfortunately, the image was created a while ago. The Windows XP image uses release-to-manufacture (RTM) code, and the Windows 2000 image has Service Pack 2 applied. So your first task is to update systems to the latest service pack level, so that the Automatic Updates client, as well as all patches and fixes, have been installed on the computers.

Note In this hands-on scenario, you may test the results using a second computer. To do so, join the computer to the domain and move its computer account to the Desktops OU. If the computer is running Windows 2000, modify the service pack deployment exercise accordingly.

Exercise 1: Download and Extract the Service Pack

1. Create a folder on the C drive and name the folder ServicePack.
2. From the Microsoft download site, *http://www.microsoft.com/downloads*, or from the Windows XP site, *http://www.microsoft.com/windowsxp*, download the latest service pack. Save it to the C:\ServicePack folder.
3. Open a command prompt and type **cd C:\ServicePack** to change to the ServicePack folder.
4. Type ***xpsp1.exe* -x**. Substitute *xpsp1* with the filename for the most recent service pack you downloaded.
5. You will be prompted to indicate the location to which the service pack will be extracted. Type **C:\ServicePack**.
6. The service pack is extracted. Use Windows Explorer to navigate the folder structure that was created. Make note of the location of update.exe (in the Update folder), which is used to launch installation of the service pack on a single machine, and of update.msi (in the same folder), which will be used to deploy the service pack through group policy-based software distribution.

Exercise 2: Deploy the Service Pack with Group Policy

1. Share the C:\ServicePack folder with the share name ServicePack.
2. Open Active Directory Users And Computers.
3. Expand the domain and locate (or create) the Desktops OU.
4. Create a computer object in the Desktops OU called Desktop0569 to represent one of the new systems.

Note If you have a second system with which to perform this case scenario exercise, move that system's account into the Desktops OU.

5. Right-click the Desktops OU and choose Properties.
6. Click the Group Policy tab.
7. Click New to create a new GPO. Name the object SP-Deploy.
8. Select the SP-Deploy group policy link and click Edit. Group Policy Object Editor opens.
9. Navigate to Computer Configuration\Software Settings.
10. Right-click Software Installation and choose New, then Package.
11. Type the path **\\server01.contoso.com\servicepack** and press Enter. The browse dialog box will take you to the root of the extracted service pack.
12. Navigate to the Update.msi file you identified in the previous exercise. Select the Update.msi file and click Open.
13. Select Assigned and click OK. The package is created.
14. Close Group Policy Object Editor and the Desktop OU's Properties dialog box.
15. (Optional) If you have a second system, you can test the deployment of the service pack. Remember that Windows XP computers are configured by default to optimize logon, so it may take two restarts before the service pack is applied. You can confirm the service pack level on a machine by clicking Start, Run, and then typing **winver**.

Exercise 3: Install SUS

1. Navigate to *http://go.microsoft.com/fwlink/?LinkId=6930*. You will be prompted to add the site to your trusted sites list, which you should do.
2. Download the SUS installation package.
3. Start SUS installation by double-clicking the downloaded file.
4. On the Welcome screen, click Next.

5. Read and accept the End User License Agreement, and then click Next.
6. Choose a Custom installation and then click Next.
7. On the Choose File Locations page, choose Save The Updates To This Local Folder. The default, C:\SUS\Content, is fine. Click Next.

Note The updates might consist of several hundred megabytes of files. If you have a slow Internet connection, or if you want to save time during this exercise, choose the second option, Keep The Updates On A Microsoft Windows Update Server instead.

8. For Language Settings, choose English Only, then click Next.
9. On the following page, choose I Will Manually Approve New Versions Of Approved Updates, then click Next.
10. The following page should indicate that the client download location is *http://SERVER01*. Click Install.
11. When installation is complete, click Finish. Internet Explorer will be opened, and will take you to the SUS administration page. Continue with Exercise 4.

Exercise 4: Synchronize SUS

1. If you are not already viewing the SUS administration page, open Internet Explorer and navigate to *http://SERVER01/SUSAdmin*.

Note To view the SUS administration site, you may need to add Server01 to the Local Intranet trusted site list to access the site. Open Internet Explorer and choose Internet Options from the Tools menu. Click the Security Tab. Select Trusted Sites and click Sites. Add **Server01** and **Server01.contoso.com** to the trusted site list.

2. Click Synchronize Server on the left navigation bar.
3. Click Synchronization Schedule.

   You will manually synchronize for this exercise. However, you can examine synchronization options by clicking Synchronize Using This Schedule. When you are finished exploring settings, click Cancel.
4. On the Synchronize Server page, click Synchronize Now. If you have elected to download updates to the server, synchronization may take some time.
5. After synchronization has occurred, you will be redirected automatically to the Approve Updates page. You can also click Approve Updates on the left navigation bar.

6. Approve a small number of updates so that you can return later to experiment further with approval and automatic updates.
7. Examine other pages of the SUS administration site. After you have familiarized yourself with the site, close Internet Explorer.

Exercise 5: Configure Automatic Updates

1. Open Active Directory Sites and Services.

Note Most enterprises have found little reason to link GPOs to sites, rather than OUs or the domain. However, SUS-related policies lend themselves well to site application, since you are directing clients to the most site-appropriate SUS server.

2. Right-click the Default-First-Site-Name site and choose Properties.
3. Click the Group Policy tab.
4. Click New and name the new GPO **SUS-Site1**.
5. Click Edit. The Group Policy Object Editor opens.
6. Navigate to Computer Configuration\Administrative Templates\Windows Components\Windows Update.
7. Double-click the policy: Specify Intranet Microsoft Update Service Location.
8. Click Enabled.
9. In *both* text boxes, type **http://server01.contoso.com**.
10. Click OK.
11. Double-click the policy: Configure Automatic Updates.
12. Click Enabled.
13. In the Configure Automatic Updating drop-down list, choose 4-Auto Download And Schedule The Install.
14. Confirm the installation schedule: daily at 3:00 A.M.
15. Click OK.
16. Double-click the policy: Reschedule Automatic Updates Scheduled Installations.
17. Click Enabled.
18. In the Wait After System Startup (Minutes) box, type **1**.

Exam Tip The Wait After System Startup policy is used to reschedule a scheduled installation that was missed, typically when a machine was turned off at the scheduled date and time.

19. Click OK.
20. Close the Group Policy Object Editor and the Properties dialog box for Default-First-Site-Name.
21. To confirm the configuration, you can restart the server, which is also within the scope of the new policy. Open System from Control Panel and click the Automatic Updates tab. You will see that configuration options are disabled, as they are now being determined by policy.

Chapter Summary

- The Microsoft Software Update Services enable you to centralize and manage the approval and distribution of Windows critical updates and Windows security rollups. One or more SUS servers host lists of approved updates and, optionally but typically, the update files themselves. Automatic Updates clients are configured, usually through GPOs, to obtain updates from intranet SUS servers, rather than from Microsoft Windows Update.
- Software Update Services does not distribute service packs.
- Service packs can be obtained free of charge from Microsoft. If the service pack is a single file, it can be extracted from the command prompt by entering the service pack's filename followed by the -x switch.
- Service packs are deployed easily by assigning a software installation package to the computer configuration's software settings policies in a GPO.
- Tracking and managing licenses and compliance is an important part of an administrator's job. Windows Server 2003 gives you the ability to assign licenses based on concurrent connections to a specific server or to maintain a license for each device or user that connects to any number of servers in your enterprise.
- Licenses are replicated between servers' License Logging service and the site license server. The site license server can be identified using Active Directory Sites And Services, but site licensing is administered using the Licensing tool in the Administrative Tools programs group.
- A license group enables users to share one or more devices. The number of Windows Device CALs is assigned to the license group.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Have an understanding of SUS installation and configuration. Although the exam objectives will not address SUS setup directly, the way you configure SUS impacts the tasks you will perform to maintain the SUS infrastructure, so it is important to be comfortable with the "big picture" of SUS.
- Focus on SUS administrative tasks, such as synchronizing, approving updates, viewing logs and events, and configuring Automatic Updates through System in Control Panel (on a stand-alone computer) or using Group Policy in a larger environment. Remember that you cannot direct a computer to an SUS server using the Automatic Updates properties on a client. You must use Group Policy, or a registry entry, to redirect the client to an intranet server rather than Microsoft Windows Update.
- Be able to calculate license requirements in a variety of Per Server or Per Device or Per User scenarios. Remember that license groups allow multiple users to share one or more devices.

Key Terms

Client Access License The license that allows a user or device to connect to a server product for any functionality, including file and print service or authentication.

Per Server license mode Licenses are allocated when a user or device connects to the server or product. When the user disconnects, the license is returned to the available license pool. This mode requires sufficient licenses to support the maximum number of concurrent connections on each individual server.

Per Device or Per User mode Licenses requirements allow a single CAL to authorize a user (who may use more than one device) or a device (which may be used by more than one user) to connect to any number of servers.

License group Because the License Logging service allocates licenses based on user name and not device name, Windows Device CALs are given to a license group. A license group has one or more users, and is allocated licenses equivalent to the number of devices used by that group to connect to server products.

10 Managing Hardware Devices and Drivers

Exam Objectives in this Chapter:

- Install and configure server hardware devices
  - Configure driver signing options
  - Configure resource settings for a device
  - Configure device properties and settings
- Monitor server hardware. Tools might include Device Manager, the Hardware Troubleshooting Wizard, and appropriate Control Panel items.

Why This Chapter Matters

Hardware devices give us access to the information that is processed by a computer. From the monitor display to the keyboard, from mouse to multimedia, a computer without devices is virtually useless. Microsoft Windows Server 2003 categorizes devices, listing them for examination and configuration in Device Manager.

Devices and the software that the operating system uses to communicate with the device (a "driver") is a combination unique to the device being used and the version of the operating system on which it runs, and most drivers are not interchangeable between operating systems; that is, you cannot use a driver that is designed for use with the Windows 98 operating system with Windows XP or Windows Server 2003.

In addition to using the proper driver initially, ongoing maintenance of devices and driver configurations is necessary. Updates to drivers are common, as functional changes in operating systems and devices dictate corresponding changes in drivers. Changes must be installed in the form of a new driver, which is usually provided by the device vendor. Service Packs may also contain updated driver releases.

Given the precise match required among the device, driver, and operating system, faulty configurations quickly make for non-working devices. The correct driver for your device and the operating system for which it's designed should work properly, and as an administrator you can easily update these drivers through Device Manager.

In most environments, it is not desirable for an end-user to have the ability to install new drivers. However, administrators may not want (nor have the time) to work on each computer individually to configure all devices and their drivers properly. The configuration of driver signing options and the granting of selective privileges to the appropriate users will give the most flexibility for device configuration and installation of drivers.

Lessons in this Chapter:

Before You Begin

This chapter assumes that you have a fair, working knowledge of the most common types of computer devices such as printers, mouse devices, keyboards, network cards, and so on. Physical optimization, testing, and troubleshooting of physical devices is beyond the scope of this chapter.

Examples and practices involving the installation, configuration, and troubleshooting of devices and drivers will be performed in a Windows Server 2003 environment with standard devices. To emulate the exercises, you should have a Windows Server 2003 named Server01 as a domain controller in the *contoso.com* domain.

Lesson 1: Installing Hardware Devices and Drivers

Hardware devices communicate with the Windows Server 2003 operating system by means of a software driver. Devices and their drivers, if not installed automatically through Plug and Play, can be configured through the Device Manager.

After this lesson, you will be able to

- Understand the relationship between devices and drivers
- Use Device Manager to analyze detected devices
- Use Device Manager to install a device

Estimated lesson time: 20 minutes

Devices and Drivers

The easiest way to think about devices and their associated drivers is to divide the devices into two logical categories: Plug-and-Play (PnP) and non-Plug-and-Play (down-level) devices. Most devices manufactured since 1997 are PnP devices, and most PnP drivers for devices are included on the Windows Server 2003 installation CD. When a device is initially detected by Windows Server 2003, and if an acceptable driver is found for that device, the device will be installed and such resources as interrupt requests (IRQs) and direct memory access (DMA) will be allocated for use by the device. The device will then be listed in the categorized listing of devices in Device Manager.

If the PnP driver is not on the Windows Server 2003 Installation CD, you will need the vendor-supplied drivers available when the Windows Server 2003 initially detects, identifies and attempts to install the device. For devices that Windows Server 2003 can identify, you will be prompted for a driver. If the request for the driver is bypassed, Windows Server 2003 will indicate the identified, non-configured device with a yellow warning icon in Device Manager. This icon, as shown in Figure 10-1, is also used if there are duplicate devices on the system or if there are conflicts between the resource demands of drivers, which is extremely rare for newer computer systems and devices.

Figure 10-1 Device Manager warning icon

If a device cannot be identified by Windows Server 2003, no request for a driver will be issued, and the unknown device will be identified with a yellow question mark in Device Manager. For a non-configured or non-identified device, you must install the appropriate driver manually for the device to function properly.

Using Device Manager

Device Manager provides a view, similar to Windows Explorer, of the hardware that is installed on your computer. You can use Device Manager to update the drivers for hardware devices and modify settings related to devices. Device Manager is accessible through the Control Panel by selecting System, the Hardware tab on the Systems Properties dialog box, and then Device Manager to access the Device Manager Properties page, or as part of the Computer Management console, accessible from Administrative Tools. Table 10-1 describes the tasks for which Device Manager can be used.

Table 10-1 Device Manager Tasks

Task	Usage
Determine whether the hardware on your computer is working properly	Properly configured devices are listed by category. Detected devices that are not configurable, either because of a lack of an appropriate driver or an irresolvable resource conflict, are indicated by a yellow icon with an exclamation point. Devices that cannot be identified are indicated by a yellow question mark icon.
Print a summary of the devices that are installed on your computer	On the Action menu in Device Manager, select Print. Print options include System Summary, Selected Class or Device, and All Devices And System Summary.
Change hardware configuration settings	Right-clicking and choosing Properties (or double clicking) on any device will open the Properties page for the device.

Table 10-1 Device Manager Tasks (Continued)

Task	Usage
Device Properties Pages	
General tab	Identify the device type, manufacturer, location, and status of the device. The device can also be enabled or disabled from the Device Usage drop-down list.
Driver tab	View details of the device driver such as driver version, driver provider, and whether the driver has been digitally signed; install updated device drivers; update the device driver; roll back to a previously installed version of the driver.
Resources tab	Lists the resource usage by a device, including I/O ranges, memory addresses, and IRQ use. The ability to disable automatic configuration, which enables manual configuration, varies by device: Some devices do not allow for manual configuration of resources.

Exam Tip You can use Device Manager to manage devices only on a local computer. On a remote computer, Device Manager will work only in read-only mode.

A list of devices, drivers, and system configuration can be printed through the Print command on the Action menu in Device Manager or output to a comma-separated-values (CSV) file using the Driverquery command-line utility, the parameters for which are listed in Table 10-2.

Table 10-2 Driverquery Command Parameters

Parameter	Output
/S system	Specifies the name or Internet Protocol (IP) address of a remote computer to connect to. The default is the local computer.
/U domain\user	Runs the command within the context of the user specified by User or Domain\User. The default is the permissions of the user who is logged on to the computer issuing the command.
/P password	Specifies the password of the user account that is specified in the /U parameter.
/FO format {TABLE | LIST | CSV}	Specifies the format to display the driver information. Valid values are TABLE, LIST, and CSV. The default format for output is TABLE.
/NH	Omits the header row from the displayed driver information. Valid when the /FO parameter is set to TABLE or CSV.
/V	Specifies that detailed driver information be displayed. Not a valid option for signed drivers.

Table 10-2 Driverquery Command Parameters (Continued)

Parameter	Output
/SI	Specifies to display the properties of signed drivers.
/?	Displays help at the command prompt.

Users, Administrators, and Device Installation

As with most installation tasks, administrators have the ability to install any device and its associated drivers. Users, on the other hand, have very limited ability to install devices on the computer. By default, users can install only PnP devices, with the following considerations:

- The device driver has a digital signature.
- No further action is required to install the device, requiring Windows to display a user interface.
- The device driver is already on the computer.

If any of these conditions is not met, the user cannot install the device unless delegated additional administrative authority.

Exam Tip If a PnP device requires no additional user interaction for installation, and the driver is already on the computer, a default user can connect and use the device. This applies to any universal serial bus (USB), parallel, IEEE 1394 device, especially printers. The Load And Unload Device Drivers user right, configurable through Group Policies, does not apply to PnP drivers, and need not be enabled for a user to install a PnP device.

Driver Signing Options

Device drivers and operating system files included with Windows 2000 or higher have a Microsoft digital signature. The *digital signature* indicates that a particular driver or file was not altered or overwritten by another program's installation process. Device drivers provided by vendors outside of Windows 2000 or higher may or may not be signed.

You can control how the computer responds to these unsigned driver files during their installation. These settings are configurable through Control Panel by selecting System, the Hardware tab on the Systems Properties dialog box, and then Driver Signing to access the Driver Signing Options Properties page on an individual computer. The options for unsigned driver installation behavior are:

- **Ignore** To allow all device drivers to be installed on the computer, regardless of whether they have a digital signature. This option is available only if you are logged on as an administrator or as a member of the Administrators group.
- **Warn** To display a warning message, allowing you to allow or deny driver installation, whenever an installation program or Windows attempts to install a device driver without a digital signature. This is the default behavior.
- **Block** To prevent an installation program or Windows from installing device drivers without a digital signature.

Group Policy is an effective tool for simultaneously changing the Driver Signing Options setting on multiple computers. To prohibit a user from changing the setting on his or her computer, you must deny access to the Hardware Properties pages in Control Panel and disable the MMC snap-in for Device Manager in the Computer Management console. These settings will not change the user's ability to install PnP devices.

Practice: Installing Device Drivers

In this practice, you will install a network adapter, change the Driver Signing Options, and then return the computer to its default configuration.

Exercise 1: Install a Network Adapter

1. Open the System Properties page from Control Panel, and then on the Hardware tab, click Add Hardware Wizard.
2. Click Next and wait for the Hardware Wizard to scan your computer for new devices. If you have not added any devices, the wizard will ask whether the new device has been connected.
3. Select Yes, I Have Already Connected The Hardware, and then click Next.
4. From the Installed Hardware list, scroll to the bottom, select Add A New Hardware Device, and then click Next.
5. Select the Install The Hardware That I Manually Select From A List (Advanced) option, and then click Next.
6. From the Common Hardware Types list, select Network Adapters, and then click Next.
7. Select Microsoft as the Manufacturer, and Microsoft Loopback Adapter as the Network Adapter, and then click Next.
8. Click Next, and then Finish, to close the wizard.

Windows Server 2003 will now load the driver and install the device. The network adapter named Microsoft Loopback Adapter will appear in Device Manager under the Network Adapters category.

Exercise 2: Set Driver Signing Options

1. Open the System Properties page from Control Panel, and then on the Hardware tab, click Add Driver Signing.
2. Select the Block option.
3. Click OK.

You have now disallowed the installation of unsigned drivers.

Exercise 3: Return Computer to Default

1. Open Device Manager. Right-click Microsoft Loopback Adapter and choose Uninstall from the shortcut menu.
2. Click OK to confirm the device's removal.
3. Close Device Manager.
4. Open the Driver Signing Properties page again, and select Warn.
5. Select Make This Action The System Default.
6. Click OK twice.

You have returned your computer to its default configuration.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You want to make certain that no unsigned drivers are used on the desktop computers in your environment. What Driver Signing settings and related configuration will assure this condition?

   _____

   _____

2. A user wants to install a USB printer connected to their his or her computer. The drivers for the printer are included with Windows Server 2003. Can the user install the printer?

   _____

   _____

3. A user wants to install a USB printer connected to his or her local computer. The driver is provided by the vendor, and is not included with Windows Server 2003. The driver is digitally signed. Can the user install the printer?

Lesson Summary

- Device Manager lists all detected devices, and indicates problems with identification or driver configuration.
- Driver configuration can be output to a printed document using Device Manager, or to a CSV file using the Driverquery command.
- Users can connect and install any completely PnP device. If any user intervention is required, a user will not be able to install a device.
- Interface access points to device and driver configuration can be disabled through local and domain-based Group Policies.
- Unsigned Driver Installation behavior has three settings: Ignore, Warn, and Block.

Lesson 2: Configuring Hardware Devices and Drivers

Devices may require updated drivers due to changes in the Windows Server 2003 operating system or changes in the way that a vendor programs a device to function. Drivers can be updated through Device Manager.

To minimize the impact of possible problems with a new driver, a feature of Device Manager allows for a return to the previous driver. This rollback feature is accessible through the Properties page of the device.

Occasionally, the automatic resource configuration within Windows Server 2003 is insufficient to accommodate a unique pattern of device use on a particular computer. If a device needs to have static resources (IRQ, I/O Port, DMA, or Memory Range) set, Device Manager can be used to remove the Automatic Settings use in favor of a setting configured by the user/administrator.

After this lesson, you will be able to

- Use Device Manager to update, roll back, and uninstall drivers
- Use Device Manager to analyze and configure resource use by devices

Estimated lesson time: 15 minutes

Updating Drivers

In Device Manager, most devices can have their drivers updated. The driver update process is a manual one, whether the device is PnP or not, and must be accomplished by an administrator—assuming that the user has not been granted elevated privilege to do so—at the console of the local computer.

Note An exception to the requirement for local installation with administrative credentials exists if the driver is provided through Windows Update. See Chapter 9, "Maintaining the Operating System," for more information about the Software Update Services (SUS) and Windows Updates.

The process to update a driver is nearly the same as for a device that has been detected properly, but whose driver was not available at installation. After initiating the driver update process for a device from within Device Manager, the Add Hardware Wizard asks for the new driver's location and the driver is installed. Some core system drivers will require a restart of the computer after installation, but most peripheral devices will not. The Properties page where the update of a driver is started is shown in Figure 10-2.

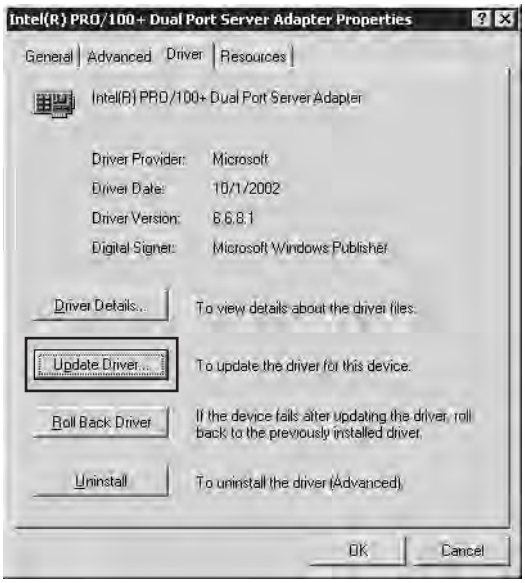

Figure 10-2 Driver update

Note If you choose to uninstall a device that was configured through PnP, you must scan for hardware changes in Device Manager to have the device reinstalled because Windows Server 2003 removes the device from the configuration even if the device is still connected to the computer.

Rolling Back Drivers

Occasionally, a new driver will not function properly and cannot be kept in the configuration for the device. If the replaced driver was performing properly, then rolling back to the previous driver can be accomplished through Device Manager. Windows Server 2003 automatically backs up the driver that is being replaced through the update driver process, making it available through the Roll Back Driver option. The Properties page where the rollback of a driver can be initiated is shown in Figure 10-3. The contrast between this feature and the Last Known Good Configuration option is discussed in the next lesson.

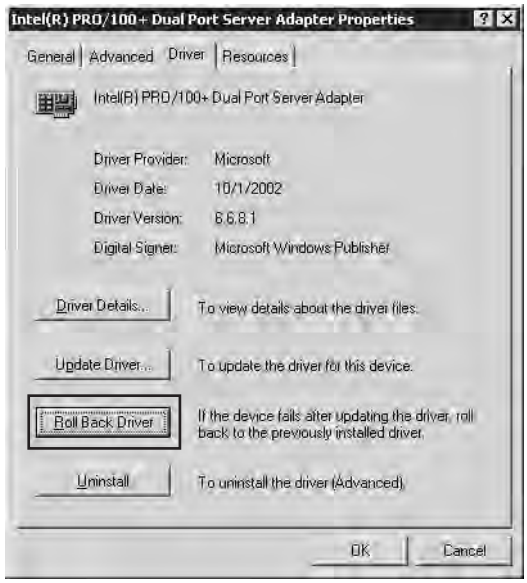

Figure 10-3 The Roll Back Driver option

Uninstalling Drivers

Drivers may be uninstalled using Device Manager. The Uninstall Driver process is initiated from the Properties page, as shown in Figure 10-4.

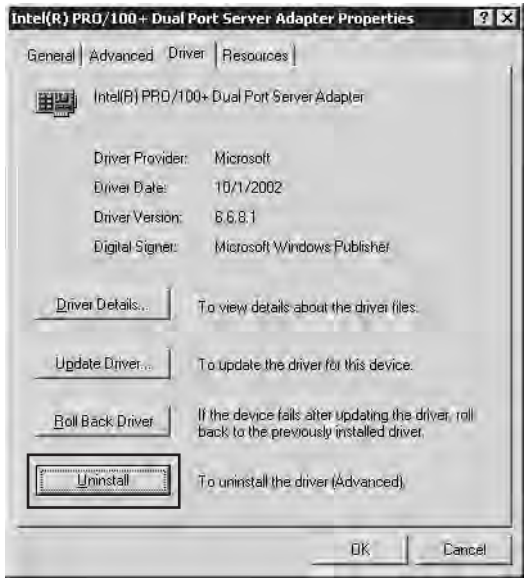

Figure 10-4 Uninstall Driver

Uninstalling a driver has different effects depending on whether the device was detected and configured through the PnP process. If the device was configured through PnP, then removal of the driver will result in the removal of the device from Device Manager as well. If the driver for the device was added manually, the device will remain in Device Manager, but will not be configured with a driver.

Resource Configuration

Devices and their drivers require system resources to communicate with and process data through the operating system. These resources are configured automatically by Windows Server 2003, sometimes in a shared capacity with other devices within the system. In circumstances where resources must be statically configured, Device Manager allows for some control of the resources assigned for use by a device. If configuration is not available, the resources used by a device and its driver cannot be configured manually. The Resources tab of a device's Properties page of a manually configurable resource is shown in Figure 10-5.

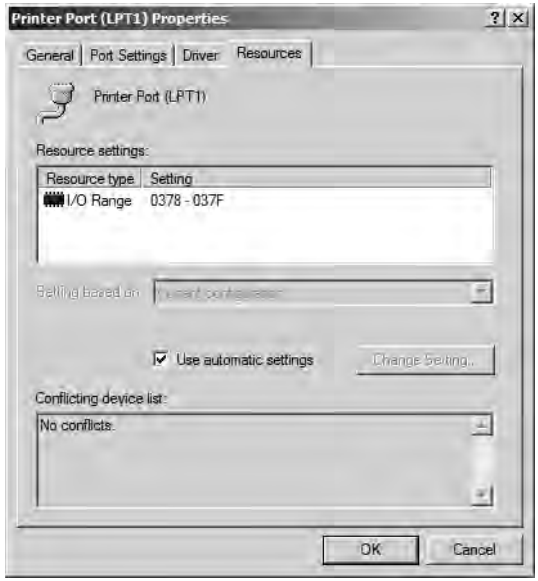

Figure 10-5 The Resources tab of a device's configurable Properties

To configure a resource assignment manually, the Use Automatic Settings check box must first be cleared, then the resources can be set.

Caution Any resources set manually make both the resource and device unavailable for automatic configuration, limiting the ability of Windows Server 2003 to make adjustments. This may cause problems with other devices.

Control Panel and Device Configuration

Several devices have Control Panel applications associated with them that allow configuration of hardware devices. The same Device Manager limitations, which are based on user rights, for the installation, updating, or removal of device drivers exist within the Control Panel applications.

Such Properties pages are administered separately through Group Policies, and can be removed from user view and access. This setting is in the User Configuration section of a Group Policy.

Practice: Configuring Devices

In the following practice, you will temporarily change the configuration of a network card to remove it from service without uninstalling the device.

Exercise 1: Disable a Device

1. Open Device Manager, then select a network card configured for your computer.
2. In Device Manager, double-click the listing of the network card.
3. Select the Device Usage drop-down list and then select Do Not Use This Device (Disable).

   The device is now disabled from operation within this Hardware profile.
4. Open the Properties page for the network card, and choose Use This Device (Enable) to re-enable the network card for use in this Hardware profile. Alternatively, you can right-click the device and select Enable or Disable, depending on the current state of the device.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Under what circumstances would you adjust the resource settings for a device?

   _____

   _____

2. You need to remove a PnP device from a configuration temporarily, but want to leave it physically connected to the computer. You want to minimize the amount of work required to use the device later. Which of the following is the best option to accomplish your goal?

   a. From the Properties page of the device, choose Do Not Use this Device (Disable).

   b. From the Properties page of the device, choose Uninstall.

   c. Using the Safely Remove Hardware utility, choose to remove the device.

3. Greg's computer has an external USB Hard Disk connected to a USB hub on his computer. He is reporting that the disk is connected properly, but the drive (G) normally associated with the disk is not available. Upon investigation, you discover that the indicator light on the hub is not illuminated and the device does not appear in Device Manager. Disconnecting and re-connecting the device has no effect. What is likely the quickest way to return the disk to proper functionality?

Lesson Summary

- Device Manager can be used to Disable/Enable individual devices.
- Manual resource configuration is possible for some devices, but should be done only when there is a conflict with other resources on the computer. Manual configuration should be kept to a minimum so as to allow Windows Server 2003 the greatest amount of flexibility in automatically configuring resources for all devices.
- Driver Updating is done through Device Manager.
- Driver Roll Back is done through Device Manager, and allows for use of a driver that was previously configured for a device.
- Uninstalling a PnP device requires rescanning of the computer to re-enable the device. Uninstalling a non-PnP device requires reinstallation to enable the device.

Lesson 3: Troubleshooting Hardware Devices and Drivers

Problems with drivers will arise, particularly when driver configuration is not possible through PnP means, or when core system component drivers are updated. When a device configuration is not possible through strictly PnP means, the chance of mismatching devices and their drivers increases. With core system component driver updates, which require a computer restart, any problems with the driver will not be known until the computer restarts.

After this lesson, you will be able to

- Understand how to use Disaster Recovery Methods for Devices
- Understand and analyze driver-related problems

Estimated lesson time: 15 minutes

Recovering from Device Disaster

Occasionally, when you install or upgrade a driver for a device, there is a problem with the functioning of that device on your system. Depending on the importance of the device, the effect of the problem will range from annoying to catastrophic. Particularly for such core system components as video drivers, a faulty configuration can render the computer unusable. Rolling back the driver, after all, is difficult if you cannot see the screen.

Thankfully, there are multiple methods of recovery from faulty driver configuration. The tools available are specifically suited to different purposes, and have varying chances of success. Tools that can be used in the event of incorrect driver configuration are listed in Table 10-3.

Table 10-3 Driver Recovery Tools

Tool	Severity	Use
Driver Rollback (Device Manager)	Low. Most system functions remain intact.	Use the Property page for the device to go back to the last driver that was working properly. Contact the vendor to resolve the issue with the new driver.
Last Known Good Configuration	Medium/High. The device driver update requires a restart, and the computer will not resume to the point of allowing you to log on.	When you change drivers that require a restart, the Registry Key HKLM\System\CurrentControlSet can be restored with the old driver information. By pressing F8 as the system restarts, you can select the Last Known Good Configuration, which restores the key. If the problem does not surface until you have successfully logged on (which is often the case with an updated video driver), Last Known Good will be of little use because it is overwritten upon successful logon.

Table 10-3 Driver Recovery Tools (Continued)

Tool	Severity	Use
Safe mode	Medium/High. System is unusable.	By pressing F8 as the system restarts, you can select Safe mode as a boot option. This mode uses only minimal system and device drivers—enough to start the computer and log on—which allows you to access Device Manager and disable the offending device.
Recovery Console	High. Last Known Good and Safe modes do not work.	The Recovery Console allows you to log on and access limited parts of the file system from a command prompt. From the Recovery Console, you can disable the device driver that is causing the problem, but you must know the correct name of the device or driver (or both), which can be cryptic.

Device Manager Status Codes

When a device fails, an error message is usually reported in Device Manager with an exclamation point in a yellow icon next to the device. If you double-click the device (or right-click the device and then click Properties), a dialog box is displayed and any error messages that Device Manager detects are listed. This Device Status has some friendly text with it, but troubleshooting may require that you understand more than the text message delivers. Often, there is a code listed with the text that gives a better idea of how to troubleshoot the problem. These codes and suggested troubleshooting strategies are listed in Table 10-4.

Table 10-4 Device Failure Troubleshooting

Code	Friendly Text	Troubleshooting Strategy
1	This device is not configured correctly. To update the drivers for this device, click Update Driver. If that doesn't work, see your hardware documentation for more information.	Use Update Driver to update the driver.
3	The driver for this device might be corrupted, or your system may be running low on memory or other resources.	The driver may be corrupted. If you attempt to load a file that is corrupted the system may think that it needs more memory. Use Task Manager to confirm that your system is not low on memory.

Table 10-4 Device Failure Troubleshooting (Continued)

Code	Friendly Text	Troubleshooting Strategy
10	The device cannot start. Try updating the device drivers for this device.	Run the Hardware Update Wizard using the Update Driver button, but do not let Windows Server 2003 automatically detect devices. Instead, select Install From A List Or Specific Location (Advanced), and manually point the wizard to the appropriate driver.
12	This device cannot find enough free resources that it can use. If you want to use this device, you will need to disable one of the other devices on this system.	Click the Resources tab on the Properties page containing the error. Windows Server 2003 will, likely, be able to enumerate the associated device that is in conflict with the device in question. Either disable or remove the device that is in conflict. You can then add the device you removed back into the system and see if the device can take new resources on its own, or if you will have to assign resources manually.
Most other codes	Various	Most other codes involve an inappropriate driver, which should be reinstalled.

Tip Remember, if a driver is signed, it is verified to work with Windows Server 2003. You can get a list of signed drivers under Software Environment of the System Information utility. System Information is accessible through the System Tools program group, or by typing winmsd at the Run line.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You have finished configuring a new display driver, and are prompted to restart the computer for the changes to take effect. Shortly after logging on, the computer screen goes blank, making working on the computer impossible. Which troubleshooting techniques or tools will allow you to recover most easily from the problem with the display driver?

   a. Last Known Good Configuration

   b. Driver Rollback

   c. Safe mode

   d. Recovery Console

2. In Device Manager, you have a device that displays an error icon. On the Properties page for the device, you read that the Device Status is: "device could not start." What course of action will solve the problem?

3. The vendor for a wireless network card installed in your computer has released a new driver. You want to test the driver for proper functionality. Which Device Manager option will you select to test the new driver?

Lesson Summary

- The Last Known Good Configuration option is useful for reverting to a previously used, non-PnP driver, but only if you have not logged on to the system after restarting.
- Starting the computer in Safe mode loads a minimal set of drivers, allowing for access to Device Manager to either disable, uninstall, or roll back a driver that is prohibiting the system from functioning properly.
- Most driver problems occur during manual configuration of an inappropriate driver.
- Resource settings should only be adjusted manually when conflicting settings cannot be resolved by the operating system.
- All manually configured resource allocations must be unique.

Case Scenario Exercise

If a computer is experiencing hardware resource allocation conflicts, hardware profiles allow for the selection of devices to be enabled in different circumstances. As an alternative to manually attempting to configure which device should be assigned what resource, and perhaps never determining a working configuration, defining a hardware profile in which a device is not enabled allows for resources to be used for other devices.

Hardware profiles also allow for the optimization of performance and some control of power usage through the disabling of devices and services that are not used in a particular situation. A laptop computer, for example, can have its battery life extended through the creation of a "mobile" profile, which disables devices that are not needed when the computer is disconnected from the network.

In this exercise, you will disable the network card for use in a hardware profile on a laptop computer.

1. On the Hardware tab of System Properties, click Hardware Profiles.
2. Copy the current profile to a new profile. Name the profile "mobile" and leave the Hardware Profiles Selection setting at the default (selects the first profile in the list if a selection is not made within 30 seconds).
3. Restart the computer. When prompted for selection of a Hardware profile, choose Mobile as the hardware profile for the system to use.
4. Log on, and open Device Manager from the Hardware tab in System Properties.
5. Right-click the network card reported in Device Manager and choose Properties.
6. In the Device Usage drop-down list on the Properties page for the network card, select Do Not Use This Device In The Current Hardware Profile (disable).

You have now disabled the network card for use in a single profile. This technique can be used in many different situations, including troubleshooting devices, by creating Hardware profiles that enable or disable different devices whose combined interactions and resource usage you are testing.

Troubleshooting Lab

The distribution files for Windows Server 2003 include most of the drivers needed to configure the latest hardware devices, and misconfiguration is very rare. For configuration conflicts that must be resolved manually, however, misconfiguration is a more common occurrence.

When a device configuration change causes the computer to fail on restart, the Last Known Good Configuration allows for rollback to use of a driver that was last in use. Assuming that logon has not been accomplished since the problematic device driver was installed, the Last Known Good Configuration is a usable option.

If logon is accomplished, the Last Known Good Configuration is overwritten with the current configuration. If a driver fails, making the computer unusable after logon, then Safe mode is a boot option that loads only a minimal set of drivers to allow configuration of malfunctioning devices and drivers.

In this lab, you will activate the Last Known Good Configuration and Safe mode options during the startup of your computer.

1. Restart your computer.
2. As the computer is starting up, press F8.

3. Activate the Last Known Good Configuration (last configuration that worked).

   At this point, all non-PnP drivers installed since the last restart and logon will have reverted to their previous state.

4. Restart your computer.
5. As the computer is starting up, press F8.
6. Start the computer in Safe mode.
7. Log on to the computer, then start Device Manager.

   You can now configure devices and their drivers for booting in Normal mode.

Chapter Summary

- You must have administrative privileges on a computer to install non-PnP devices and their drivers.
- Users are able to install true PnP devices. If the drivers need to be added to the computer, or any additional configuration or input is necessary during the installation, the user will not be able to install the device.
- Device Manager will indicate, with one of several types of icons, any devices that cannot be configured due to driver identification or resource conflict problems.
- The Device Manager and any Control Panel applications that configure hardware can be made unavailable to the user through Group Policies.
- Updated drivers can be rolled back to the previously used driver with the Roll Back Driver function of Device Manager.
- Devices can be disabled or enabled through Device Manager.
- PnP devices that have signed drivers on the Windows Server 2003 distribution CD will configure automatically, requiring no user intervention.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Review the use of Device Manager to install device drivers, update device drivers, roll back device drivers, and disable or enable devices in a hardware profile. Remember that Device Manager can change settings only on a local system—remote use of Device Manager is limited to read-only mode.
- Users can only install PnP devices.
- Administrative credentials are required to install non-PnP and vendor-supplied PnP drivers.
- Reinstallation of a driver is needed unless a resource conflict is being resolved.
- Resource conflicts are resolved by first clearing the Use Automatic Settings check box, then configuring the required resource settings.
- Last Known Good Configuration is only useful before a system restart /user logon cycle is complete.
- Safe mode will load a minimal set of drivers so that appropriate configuration can be made.

Key Terms

Roll Back Driver vs. Last Known Good Configuration A driver rollback requires logon, whereas a logon invalidates Last Known Good Configuration. Roll Back Driver and Last Known Good Configuration both revert to a previous configuration of a device driver.

Uninstalling vs. disabling a device Uninstalling a device will remove the device from all configurations. Depending on the type of device, a PnP detection might occur on the next system restart or Scan for Hardware changes. Configuration of the device on the next system restart or Scan for Hardware changes will treat the device as new.

Disabling a device maintains the driver as configured the next time that the device is enabled, but makes the device unavailable for use until enabled.

Safe Mode vs. Last Known Good Configuration Logging on in Safe mode loads a minimal set of drivers, but will not reset any drivers, whereas the Last Known Good Configuration will revert to the previous driver configuration.

11 Managing Microsoft Windows Server 2003 Disk Storage

Exam Objectives in this Chapter:

- Manage basic disks and dynamic disks
- Optimize server disk performance
- Implement a RAID solution
- Defragment volumes or partitions
- Monitor and optimize a server environment for application performance
- Monitor disk quotas
- Recover from server hardware failure

Why This Chapter Matters

If there's one truism about information technology, it's that no matter how much storage you have today, it will be full tomorrow. You probably remember when hard drives were measured in megabytes. Many organizations are now talking terabytes. And with all that data, and all those users needing all that information comes an enormous strain on the storage subsystems on your servers.

Large organizations are turning to storage area networks (SANs) made up of fiber-connected, fault-tolerant arrays of disk drives. But storage that is actually attached to your servers won't disappear quite yet, so you will want to make sure that you have configured server storage to provide the optimum balance of storage capacity, performance, and fault tolerance.

In this chapter, you will learn how to do just that: leverage one or more physical disks to address your storage requirements. You will learn about the storage options that Microsoft Windows Server 2003 provides, including flexible structures that make it easy to extend capacity, provide redundancy, and boost performance—usually without a restart! You'll also learn to configure and recover fault-tolerant disk sets created by Windows Server 2003's redundant array of independent disks (RAID) support. Finally, you will examine Check Disk, Disk Quotas and Disk Defragmenter, which will keep those drives working smoothly and perhaps delay the inevitable exhausting of their capacity.

Lessons in this Chapter:

Before You Begin

This chapter presents the skills and concepts related to disk storage. You are able to apply several concepts and skills using hands-on exercises that require the following configuration:

- A computer installed with Windows Server 2003, Standard Edition or Enterprise Edition.
- The server should have at least one disk drive with a minimum of 1 gigabyte (GB) of unallocated space.
- The computer should be named Server01 and should be a domain controller in the *contoso.com domain*.

Lesson 1: Understanding Disk Storage Options

Before you tackle the installation of a disk drive and the configuration of that drive, you must understand several important storage concepts. This lesson will introduce you to the concepts, technologies, features, and terminology related to disk storage in Windows Server 2003. You will learn about differences between basic and dynamic disk storage types, and the variety of logical volumes they support.

After this lesson, you will be able to

- Understand disk-storage concepts and terminology
- Distinguish between basic and dynamic storage
- Identify the strengths and limitations of basic and dynamic disks
- Identify the types of storage volumes supported on Windows Server 2003 managed disks

Estimated lesson time: 15 minutes

Physical Disks

Physical disks are the conglomeration of plastic, metal, and silicon that enable users to store enormous quantities of useless data and MP3s, and the occasional business document. Of course I'm being sarcastic here, but it is important to understand the difference between the physical disk, and its logical volume(s), which are discussed in the next paragraph. It is also helpful to remember that an advanced disk subsystem, such as hardware-based redundant array of independent disks (RAID) system, may consist of several physical disks, but its dedicated hardware controllers abstract the physical composition of the disk set so that Windows Server 2003 perceives and represents the disk system as a single physical disk.

Logical Volumes

A logical volume is the basic unit of disk storage that you configure and manage. A logical volume may include space on more than one physical disk. Logical volumes (also called logical disks in the context of performance monitoring) are physically distinct storage units, allowing the separation of different types of information, such as the operating system, applications, and user data. Logical volumes have traditionally been represented by a single drive letter.

As you dig into disk-related terminology, you will learn about partitions, logical drives, and volumes. Many resources will use all these terms interchangeably, which is possible because the technical distinctions between the terms are minuscule, and the user interface and command-line tools guide you clearly by exposing only the appropriate

type of logical volume based on the task you are performing. Don't get too hung up on the distinctions between the terms; they will become clear through experience if not through analysis.

Mounted Volumes

You noticed that we said, "Logical volumes have traditionally been represented by a single drive letter." That structure severely limited (to 26, says my kindergarten teacher) the number of volumes you could create on a system, and the flexibility with which those volumes could be used. Windows Server 2003's NTFS file system allows you to assign one or no drive letter to a volume. In addition, you can mount a volume to one or more empty folders on existing NTFS volumes. For example, you might create an empty folder Docs, on an existing volume with the drive letter X:, and mount a new 120 GB logical volume to that folder. When users navigate to X:\Docs, the disk subsystem redirects the input/output (I/O) requests to the new volume. All of this is transparent to the user.

The possibilities using this powerful feature are, as they say, "limitless." By mounting a volume to a folder path, you can extend the available drive space on an existing volume. If the existing volume is not fault-tolerant, but the new volume is fault-tolerant, the folder to which the volume is mounted, X:\Docs, represents a fault-tolerant portion of the existing volume's namespace. You could, theoretically, mount all logical volumes on a server to folders on the server's C or D drive and thereby unify enormous storage capacity under the namespace of a single drive letter.

Fault Tolerance

Fault tolerance refers to a system's ability to continue functioning when a component—in this case, a disk drive—has failed. Windows Server 2003 allows you to create two types of fault-tolerant logical volumes: mirrored (RAID-1) and striped with parity (RAID-5). You will learn more about the details of these configurations later in the chapter, but it is important to remember several facts about Windows Server 2003 fault tolerance, often called software RAID:

- In fault-tolerant disk configurations, two or more disks are used, and space is allocated to store data that will enable the system to recover in the event of a single drive failure.
- The fault tolerance options supported by Windows Server 2003 do not provide a means for a disk volume to continue functioning if two or more disks fail.
- The operating system allows you to use any two or more disk drives to create fault-tolerant volumes. You do not have to purchase any additional hardware or software to benefit immediately from fault-tolerant server configurations. However, if you use Windows Server 2003 mirrored or RAID-5 volumes, it is best practice to

use similar or identical disk drives on the same bus. Combining a variety of disk hardware, or using drives connected to a variety of small computer systems interface (SCSI) or Integrated Device Electronics (IDE) buses can affect performance significantly.

- Speaking of performance, Windows Server 2003 fault tolerance is using processor cycles and other server resources to manage the volumes. RAID-5 can be particularly detrimental to server performance. It is possible, and affordable these days, to purchase hardware-based fault-tolerant disk arrays, known as *hardware RAID*. Hardware RAID uses dedicated controllers to manage fault tolerance, and such systems are generally faster and more flexible in both management and recovery than is Windows Server 2003 RAID.
- Because hardware RAID controllers offload the management duties from the operating system, a hardware RAID array appears to Windows Server 2003 as a single disk.

Separation of Data

It is a good idea to analyze storage requirements carefully before configuring the disk subsystem of a server. Administrators typically elect to install the operating system on a logical volume separate from applications and data. By isolating the operating system, it is easier to secure the operating system volume and to manage disk space so that the volume does not run out of space. It is also usual to configure some kind of fault tolerance for the operating system.

Applications are generally stored in a separate volume, and user data and files in a third. Again, isolation of data types allows you to manage security, performance, and fault tolerance separately for each data type. If an application uses a transaction log to prepare entries into a database, as do Microsoft Active Directory directory service and Microsoft Exchange Server, it is typical to store those logs in volumes that reside on physical disks separate from the database itself, allowing the application to rebuild the database from the logs if the database fails.

Once you have thoroughly analyzed your storage requirements as they relate to the data type, security, performance, and fault tolerance, you can begin to determine how many disks you require and how those disks should be configured.

Basic and Dynamic Disks

An operating system must have a way to make sense of the physical space on a disk drive. There are two structures that Windows Server 2003 can apply to help it apportion and allocate drive space: basic and dynamic storage, also called basic and dynamic disks.

Basic Disks, Partitions, and Logical Drives

Basic disks maintain the structure with which you are probably most familiar. Each basic disk is partitioned, and each partition functions as a physically separate unit of storage. The information about the location and size of each partition is stored in the partition table of the Master Boot Record (MBR) on the drive. A basic disk can contain as many as four partitions, consisting of either four primary partitions or three primary partitions and one extended partition.

The logical volumes on a basic disk are primary partitions and logical drives. The logical volume, as mentioned, can be represented by zero or more drive letters and can be mounted to folders on an existing NTFS volume.

- **Primary partition** Each primary partition maintains one logical volume on a basic disk. If a basic disk is used to start the operating system, one and only one primary partition on the disk must also be marked as active.

Tip The computer's basic input/output system (BIOS) looks to the active partition to locate the hardware-specific files required to load the operating system. That partition is technically referred to as the *system partition* and is usually assigned drive letter "C". Once the boot process has begun, the operating system is loaded. Most servers are configured with the operating system on the C drive as well. The partition on which the operating system is stored is called the *boot partition*. Yes, it can get confusing, particularly because the same volume is referred to by the variable *%Sysvol%*. Fortunately, it's not a distinction you're likely to need to know, since most installations are completely on drive C, making the C drive the system partition, the boot partition, and *%Sysvol%*.

- **Extended partition** A basic disk may also contain an extended partition. Unlike primary partitions, extended partitions are not formatted or assigned drive letters. Instead, extended partitions are further divided into logical drives. Logical drives are logical volumes on a basic disk.

In earlier versions of Microsoft operating systems, including Windows 95, Windows 98, and MS-DOS, the operating system could only "see" the primary partition on which it was installed, plus the extended partition on the drive, if one existed. If you wanted additional storage segments on the drive, you had to configure an extended partition and apportion it into one or more logical drives. Because Windows NT, Windows 2000, Windows XP, and Windows Server 2003 can access all partitions on a disk, you only need an extended partition if you want more than four logical drives on a single disk.

Dynamic Disks and Volumes

Microsoft Windows 2000, Windows XP, and the Windows Server 2003 family also support dynamic storage. The storage units on dynamic disks are called volumes, and the first distinctions between basic and dynamic storage are that dynamic disks support an

unlimited number of volumes, and that the configuration information about the volumes is stored in a database controlled by the Logical Disk Manager (LDM) service.

The logical volume of dynamic disks is the volume. Dynamic disks support simple volumes on a single disk. When a computer has more than one dynamic disk, you are provided more storage options from which to choose. Spanned, mirrored (RAID-1), striped (RAID-0), and striped with parity (RAID-5) volumes are logical volumes that utilize space on more than one physical disk. Each volume type uses disk space differently, and is characterized by a different level of fault tolerance. The list below summarizes the volume types, though each has nuances you will learn about as the chapter progresses.

- **Simple volume** The equivalent to a basic disk partition is a dynamic disk simple volume. Simple volumes utilize space on a single physical disk, and correspond to a single logical volume. Simple volumes can be extended by appending unallocated space on other regions of the same disk, allowing you to adjust a volume's capacity with the growth of data stored in that volume. Because simple volumes exist on only one physical disk, they are not fault-tolerant.
- **Spanned volume** A spanned volume includes space on more than one physical disk. Up to 32 physical disks can participate in a spanned volume, and the amount of space used on each disk can be different. Data is written to the volume beginning with the space on the first disk in the volume. When the space on the first disk fills, the second disk is written to, and so on. Spanned volumes provide an option for increasing drive capacity. If a simple or spanned volume is filling up, you can extend the volume onto additional new storage capacity.

  But spanned volumes are not fault-tolerant, and cannot participate in any fault-tolerant configurations. Because their size tends to be greater, and because multiple physical disks are involved, the risk for failure increases. If any one disk in a spanned volume is corrupted or lost, data on the entire volume is lost as well. For these reasons, Windows Server 2003 will not allow the installation of the operating system on a spanned volume, nor can you extend or span the system volume. Spanned volumes are recommended only as a stop-gap measure when an existing volume fills to capacity, or else in situations where tolerance for failure is high—for example, a large library of read-only data that can easily be restored from tape backup in the event of failure.
- **Striped volume** A striped volume (RAID-0) combines areas of free space from multiple hard disks into one logical volume. Unlike a spanned volume, however, data is written to all physical disks in the volume at the same rate. Because multiple spindles are in use, read and write performance is increased almost geometrically as additional physical disks are added to the stripe. But like extended simple volumes and spanned volumes, if a disk in a striped volume fails, the data in the entire volume is lost.

- **Mirrored volume** A mirrored volume (also known as RAID Level 1, or RAID-1) consists of two identical copies of a simple volume, each on a separate hard disk. Mirrored volumes provide fault tolerance in the event that one physical disk fails.
- **RAID-5 volume** A RAID-5 volume is a fault-tolerant striped volume. Space on three or more physical disks is unified as a single volume. Data is written to all physical disks at the same rate, but unlike a striped volume, the data is interlaced with checksum information, called parity. Should a single disk in the volume fail, the data on that disk can be regenerated through calculations involving the remaining data and the checksum information. It is an interesting technical note that parity is distributed among all volumes in the RAID-5 set.

Basic vs. Dynamic Disks

So now that you know about basic and dynamic storage, and the types of partitions, logical drives, and volumes they support, which is better? The answer, as is frequently the case, is: "It depends."

Dynamic disks that store data are easily transferred between servers, allowing you to move a disk from a failed server to a functioning server with little downtime. Dynamic disks flex their muscle when there is more than one dynamic disk in a computer. Each Windows 2000, Windows XP, and Windows Server 2003 computer can support one disk group, which itself can contain multiple dynamic disks. The LDM database is replicated among all disks in the disk group, which increases the resiliency of disk configuration information for all the group's disks. In addition, disks can be configured to work together to create a variety of flexible and powerful volume types including spanned volumes, striped volumes (RAID-0), mirrored volumes (RAID-1), and striped-with-parity volumes (RAID-5).

Basic disks will continue to be used, however, for several reasons:

- Basic storage is the default in Windows Server 2003, so all new disks are basic disks until you convert them to dynamic—a simple process you will learn in Lesson 2.
- Dynamic disks do not offer advantages over basic disks in a computer that will have only one disk drive.
- The behavior of the LDM database also makes it difficult to transfer a dynamic disk used for starting the operating system to another computer when the original computer fails.
- Dynamic disks are not supported for removable media, and are not supported on laptops.
- Basic storage is the industry standard, so basic drives are accessible from many operating systems, including MS-DOS, all versions of Microsoft Windows, and most non-Microsoft operating systems (there are a few). Therefore, dynamic disks

cannot be used if you need to dual-boot an earlier operating system that requires access to the disks. Keep in mind that we are talking about *local* access only. When a client of any platform accesses files over the network, the underlying storage and volume type are transparent to the client.

Exam Tip Multiboot scenarios are less common these days with the advent of virtual machine technology (see *http://www.microsoft.com/windowsserver2003/techinfo/overview/virtualization.mspx*). However, if you implement a multibooted system with Windows Server 2003 as one of the operating systems, you should install each operating system on a separate, primary partition. Other configurations are risky at best. For more information on multibooting, open the Help and Support Center and search using the keyword *multiboot*.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You are installing a new 200 GB disk drive. You want to divide the disk into five logical volumes for the operating system, applications, user home directories, shared data, and a software distribution point. The drive space should be distributed equally among the five logical volumes. You also want to leave 50 GB as unallocated space for future extension of a logical volume. Considering basic and dynamic disks and the types of logical volumes they support, what are your configuration options?

   _____

   _____

2. Which of the following provides the ability to recover from the failure of a single hard drive?

   a. Primary partition

   b. Extended partition

   c. Logical drive

   d. Simple volume

   e. Spanned volume

   f. Mirrored volume

   g. Striped volume

   h. RAID-5 volume

3. You are dual-booting a system in your test lab. The computer has Windows NT 4 installed on the first primary partition, and Windows Server 2003 installed on the second primary partition. The computer is running low on disk space, so you add a new disk drive. You boot to Windows Server 2003 and configure the drive as a dynamic disk. When you later restart to Windows NT 4, you are unable to see the disk. Why?

4. To provide fault tolerance, maximum performance, and the ability to hot-swap a failed drive, you purchase a seven-disk hardware RAID array. After installing the array, you see only one new disk on Windows Server 2003. Why?

Lesson Summary

- Disk terminology can be confusing, but in the end a logical volume is almost synonymous with the terms *partition*, *logical drive*, or *volume*.
- Windows Server 2003 supports basic and dynamic disks. Basic disks support as many as four partitions: four primary partitions or three primary partitions and one extended partition, which supports multiple logical drives. Dynamic disks support simple volumes or, when more than one dynamic disk is configured, spanned, mirrored, striped, and RAID-5 volumes.
- Fault tolerance is provided by mirrored (RAID-1) volumes, which maintain a full copy of the volume's data on each of two disks, and striped-with-parity volumes (RAID-5), which stripe the data across multiple disks and use parity information to calculate data missing from any one failed disk.
- Simple volumes, spanned volumes, striped volumes (RAID-0), and all basic disk logical drives are not fault-tolerant. All data is lost if any disk supporting such volumes fails. The larger those volumes, or the more physical disks supporting those volumes, the greater the likelihood of failure.

Lesson 2: Configuring Disks and Volumes

In this lesson, you will apply the concepts of disk storage covered in Lesson 1 to the actual skills needed to install, configure, and manage disk storage. You will learn how to use the Disk Management tool to direct the detection and initialization of newly installed disks, and to apportion that disk to partitions, logical drives, and volumes. In the event that a volume fills up, you will learn how to extend that volume's capacity. And you will explore the processes involved with moving disks between servers. Finally, you will uncover the powerful new DISKPART command, which allows you to manage storage from the command line.

After this lesson, you will be able to

- Install and initialize a physical disk
- Manage the configuration of logical volumes on basic and dynamic drives
- Mount a volume to a folder on an NTFS volume
- Extend a volume's capacity
- Move disks between servers
- Convert basic and dynamic disks
- Perform disk management tasks using DISKPART

Estimated lesson time: 25 minutes

Disk Management

Disk management activities are performed using the cleverly named Disk Management snap-in, which is part of the Computer Management console. Open the Disk Management snap-in in the Computer Management console, or add the snap-in to a custom console.

Tip There is a stand-alone Disk Management console, but it is not visible in your Administrative Tools folder. Click Start, choose Run and type **diskmgmt.msc** to open the stand-alone console.

Disk Management can manage disk storage on local or remote systems. The snap-in does not manipulate disk configuration directly; rather, it works in concert with Dmadmin, the Logical Disk Manager Administrative Service that is started on the computer you are managing when you start the Disk Management snap-in.

The Disk Management interface is shown in Figure 11-1. The top frame—the list view—displays information about each partition, logical drive, or volume. The bottom frame—the graphical view—depicts disk space allocation per physical disk, as

perceived by Windows Server 2003. You can right-click the volumes in either frame to access a shortcut menu to format, delete, or assign a drive letter to the volume. If you right-click an area of unallocated disk space, you can create a partition or volume. By right-clicking the disk drive's status box, on the left of the disk's graphical view, you can initialize a new disk, convert between basic and dynamic disks, and access the disk's hardware properties dialog box.

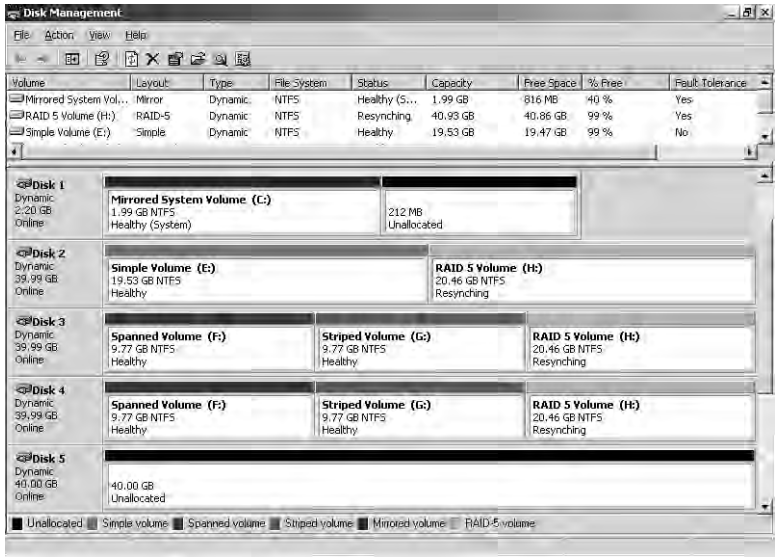

Figure 11-1 Disk Management console

Configuring Disks and Volumes

Configuring storage entails the following steps:

1. Physically installing the disk(s).
2. Initializing the disk.
3. On a basic disk, creating partitions and (if an extended partition) logical drives or, on a dynamic disk, creating volumes.
4. Formatting the volumes.
5. Assigning drive letters to the volumes, or mounting the volumes to empty folders on existing NTFS volumes.

You must be a member of the Administrators or Backup Operators group, or have been otherwise delegated authority, to perform these tasks, although only administrators can format a volume.

Installing the Disk

To add a new disk to a computer, install or attach the new physical disk (or disks). Open Disk Management and, if the drive has not been detected automatically, right click the Disk Management node and choose Rescan Disks. If a system must be taken offline to install a new disk, restart the computer, then open Disk Management. If the new disks are not automatically detected, rescan the disks.

Initializing the Disk

When you add a disk to a server, you will need to initialize that disk before you can begin to allocate its available space to partitions, logical drives, and volumes. Initializing a disk allows the operating system to write a disk signature, the end of sector marker (also called signature word), and an MBR or globally unique identifier (GUID) partition table to the disk.

If you start the Disk Management console after installing a new disk, the Initialize Disk Wizard will appear automatically. To initialize a disk manually using Disk Management, right-click the disk's status box and choose Initialize Disk.

Note On an Itanium computer, you will be prompted to select the partition style. Itanium computers containing multiple disks support two partition styles, GUID partition table (GPT) and MBR. The system partition on an Itanium computer uses the Extensible Firmware Interface (EFI) and the GPT partition style to support the 64-bit editions of the Windows Server 2003 family. More information about GPT partitions and EFI can be found in the Help And Support Center.

Creating Partitions and Volumes

After you have initialized the disk, you can begin to implement a storage structure of partitions, logical drives, or volumes.

A newly initialized disk is configured by default as a basic disk. If you wish to maintain the disk as a basic disk, you can divide the basic disk into primary and extended partitions by right-clicking unallocated space and choosing New Partition. If you choose to create a primary partition, the partition becomes a logical volume. After creating an extended partition, right-click the partition again and choose New Logical Drive. As you'll remember from earlier discussions, logical drives are logical volumes on an extended partition.

If you want to configure the disk as a dynamic disk, right-click the disk's status box in Disk Management and choose Convert To Dynamic Disk. You can then right-click the unallocated space on the disk and choose New Volume. The New Volume Wizard will

step you through the creation of supported volume types. The Select Volume Type page of the wizard is shown in Figure 11-2.

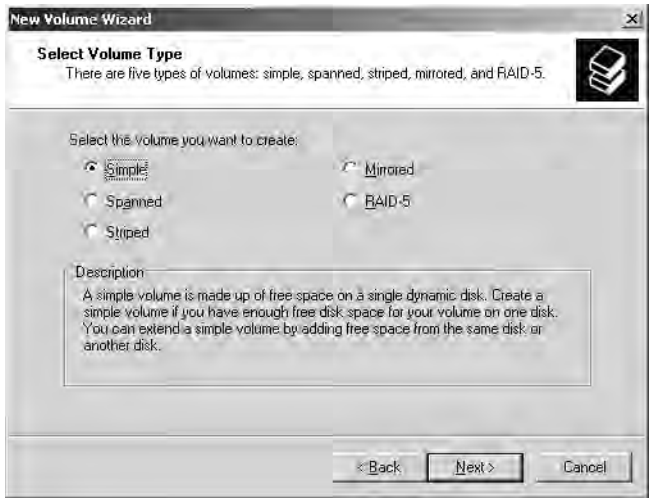

Figure 11-2 The Select Volume Type page of the New Volume Wizard

You can convert an existing basic disk to a dynamic disk—a solution that will be discussed later in this lesson.

Formatting Volumes

Windows Server 2003 supports three file systems: FAT, FAT32, and NTFS. Let's keep this discussion simple: use FAT or FAT32 only when you have very specific reasons for doing so. Only NTFS gives you the level of stability, resiliency, scalability, flexibility and security required by most organizations. Many core components of Windows Server 2003, such as file security, and services, including Active Directory and Remote Installation Services (RIS), require NTFS. All advanced storage management tasks, including multidisk volumes and disk quotas require NTFS. If you think you need FAT32, think again, then think again.

Assigning Drive Letters or Mounting Volumes

When you create a volume, it defaults to the next available drive letter. The New Volume Wizard and New Partition Wizard give you a chance to specify an alternative representation for the new logical volume. You can also right-click an existing volume and choose Change Drive Letter and Paths.

A volume can be represented by only one drive letter, though you can configure a volume to have no drive letter. However, you can mount a volume in one or more empty folders on local NTFS volumes. In the Change Drive Letter And Paths dialog box, you

can click Remove or Change to delete or modify an existing drive letter or folder mounting for the volume.

Note You cannot change the drive letter of the volume that is a system partition or boot partition.

Click Add to add a drive letter or mount point. Figure 11-3 shows a server in which the Docs folder on the X drive is a mount point to another volume. Note that the folder appears in the Explorer namespace exactly where it should, but displays a drive volume icon. When a user navigates to that folder, the user is transparently redirected to the volume.

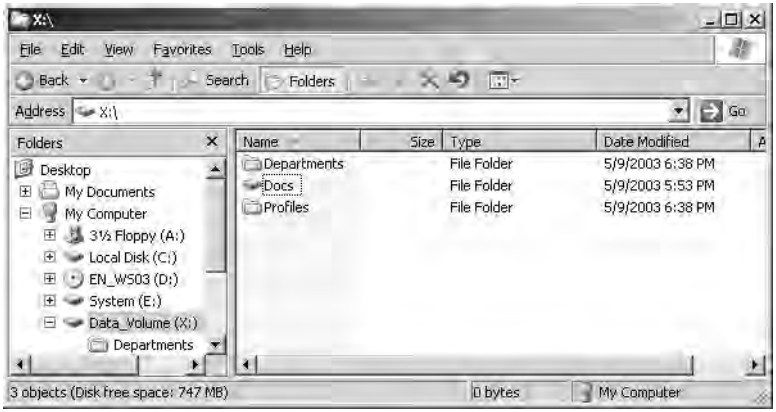

Figure 11-3 A volume mounted to a folder path

Mounting a volume in a folder on an existing volume effectively increases the target volume's size and free space. You can mount volumes regardless of whether the volumes involved are on basic or dynamic disks, and regardless of what type of volume they are. But the empty folder, the path of which becomes the path to the mounted volume, must reside on an NTFS volume. The mounted volume can, technically, be formatted as FAT or FAT32, but of course that is not the best practice.

Extending Volumes

Another way to increase a volume's capacity is to extend the volume. You can extend a simple or spanned volume on a dynamic disk so long as that volume is formatted as NTFS, and so long as the volume is not the system or boot volume. Right-click the volume and click Extend Volume. Follow the Extend Volume Wizard's instructions screen to select unallocated space on dynamic disks on which to extend the existing volume. If you extend a simple volume onto space on another physical disk, you create a spanned volume.

You can extend a partition on a basic disk using the DISKPART command. The basic partition must be formatted as NTFS, must not be the system or boot partition, and must be extended onto immediately contiguous space on the same physical disk that is either unallocated and unformatted, or formatted with NTFS.

Moving Disks Between Servers

It is possible to move disks between computers. If, for example, you plan to take a server offline, you might attach its physical disks to another server so that data can continue to be accessed. The process for doing so is the following:

1. Check the health of the disk while it is in the original server. It is recommended to open Disk Management and confirm that the disk status displays Healthy before moving the disk. If the disk is not healthy, repair the disk.
2. Uninstall the disk in the original computer. If the original server is online, uninstall the disk by right-clicking the disk in Device Manager and choosing Uninstall.
3. Remove a dynamic disk correctly. If the original server is online, open Disk Management, right-click the dynamic disk and choose Remove. This step is not necessary or possible with basic disks.
4. Physically detach the disk. If the computer supports hot-swapping the drive, you may remove the drive. Otherwise, shut down the computer to remove the physical disk.
5. Attach the disk to the target server. Open Disk Management and, if the drive has not been detected automatically, right click the Disk Management node and choose Rescan Disks. Otherwise, shut down the target server before adding the physical disk.
6. Follow instructions in the Found New Hardware Wizard. If the wizard does not appear, open Device Manager and see if the drive was detected and installed automatically. If not, open Add Hardware from Control Panel.
7. Open Disk Management. Right-click Disk Management and choose Rescan Disks.
8. Right-click any disk marked Foreign and choose Import Foreign Disks. Importing a disk reconciles the LDM databases on a new dynamic disk with the existing disks.

Some important notes about moving physical disks:

- If an imported disk contains volumes that span to other physical disks, you must attach and import all physical disks before the volumes can be accessed.
- If you move drives from several computers to a single computer, move all drives from one computer before beginning to move drives from the next computer.

- A basic volume that is moved to a new computer receives the next available drive letter. Dynamic volumes retain the drive letter they had on the original computer. If a dynamic volume did not have a drive letter on the previous computer, it does not receive a drive letter when moved to another computer. If the drive letter is already used on the computer where they are moved, the volume receives the next available drive letter.
- Use the Mountvol /n or the DISKPART automount commands to prevent new volumes from being automatically mounted and assigned a drive letter. If these commands have been used, when you add a new disk you must manually mount the volumes and assign drive letters or paths.

Converting Disk Storage

You can convert a basic disk to a dynamic disk. If the disk already contains partitions and logical drives, those units will be converted to the equivalent units for a dynamic disk: simple volumes. The structure of data on the disk is not modified, so it is possible to convert a basic disk that already contains data, although it is always best practice to back up volumes before performing disk management tasks.

To convert a basic disk to a dynamic disk, right-click the disk's status box and choose Convert To Dynamic Disk. It's that simple. If you convert a disk that contains a system or boot partition, the computer must restart.

Tip Do not convert basic disks to dynamic disks if they contain multiple operating systems (for example, the disk is set up to dual-boot with another operating system). After the disk is converted to dynamic, you can start the operating system that you used to convert the disk, but you will not be able to start the other operating systems on the disk.

Unfortunately, the reverse process is not as straightforward. Converting back to basic storage wipes out data on the drive. So you must first back up all data on the disk. Then you must delete all existing volumes on the dynamic disk before right-clicking the disk's status box in Disk Management and choosing Convert To Basic Disk. After recreating partitions and logical drives, restore the data onto the disk. Although you can convert from dynamic to basic from a technical perspective, you are actually wiping out the disk and starting over.

Performing Disk Management Tasks from the Command Prompt

Windows Server 2003 provides command-line alternatives for disk management, including the following:

- **Chkdsk** Scan a disk for errors and, optionally, attempt to correct those errors.
- **Convert** Convert a volume from FAT or FAT32 to NTFS.

- **Fsutil** Perform a variety of tasks related to managing FAT, FAT32, or NTFS volumes.
- **Mountvol** Manages mounted volumes and reparse points.

But the granddaddy of disk management command-line tools is DISKPART. Table 11-1 summarizes the DISKPART commands that achieve common disk management tasks. Diskpart can be used interactively or can call a script. To start Diskpart interactively, type **diskpart** at the command prompt. When the Diskpart command prompt (DISKPART >) appears, type **?** at any time for help. The command's built-in documentation will appear automatically when needed to help you achieve the tasks you perform. Diskpart is also well documented in the Help And Support Center.

Table 11-1 How to Complete Common Disk Management Tasks from the Command Prompt

Task	From DISKPART>	Description

List disk, partition, and volume information	`list disk`
	`list partition`
	`list volume`
		The first command lists disk information, the second command lists partition information for the current disk, and the third command lists volume and partition information for all disks.

Create a simple volume	`create volume simple size=500 disk=2`	Typed on one line, this taks creates a simple volume 500 MB in size on disk 2.
Assign a drive letter	`select volume 4` `assign letter j`	Assigns volume 4 as the J drive.
Extend a simple volume	`select volume 4` `extend size=250 disk=2`	Extends simple volume 4 (on disk 2) with an additional 250 MB on the same disk.
Create a spanned volume	`select volume 4` `extend size=250 disk=1`	Spans a simple volume 4 (on disk 2) with an additional 250 MB on disk 1.
Delete a spanned volume	`select volume 4` `delete volume`	Deletes spanned volume 4. If volume 4 was contained on disk 1 and disk 2, the space it occupied on the two disks becomes unallocated.
Create a volume mount point	`select volume 4` `assign mount=e:\Folder1`	Assigns a volume mount point to volume 4 that is accessed from E:\Folder1.
Create a striped volume	`create volume stripe size=500 disk=1,2`	Typed on one line, this task creates a striped volume which uses 500 MB on disks 1 and 2 for a total of 1 GB of storage space.
Create a mirrored volume	`create volume simple size=00 disk=1 add disk 2`	Typed on one line, this task creates a mirrored volume which uses 500 MB on disks 1 and 2 for a total of 500 MB of fault-tolerant storage space.

Table 11-1 How to Complete Common Disk Management Tasks from the Command Prompt (Continued)

Task	From DISKPART>	Description
Break a mirror	`select volume 5` `break disk 2`	Selects the mirror on volume 5 and break the mirror on disk 2.
Remove a mirror	`break disk 2 nokeep`	Deletes a mirror and remove the previously mirrored data on disk 2.
Create a RAID-5 stripe	`create volume raid` `size=500 disk=1,2,3`	Typed on one line, this task reates a RAID-5 volume which uses disks 1, 2, and 3 for a total of ~1 GB of fault-tolerant storage space.
Convert a disk from basic to dynamic storage	`select disk 2` `convert dynamic`	Converts disk 2 from basic storage to dynamic storage.
Convert a disk with unallocated space from dynamic to basic storage	`select disk 2` `convert basic`	Converts disk 2 from dynamic to basic.

Practice: Configuring Disks and Volumes

In this practice, you will use the Disk Management snap-in and Diskpart to perform a variety of disk-management tasks on Disk 0. Disk 0 must be configured as a basic disk and contain at least 1 GB of unallocated space to complete this exercise.

Exercise 1: Creating a Partition Using the Disk Management Snap-in

1. Log on to Server01 as Administrator and open the Disk Management snap-in in the Computer Management console.

   The Volume list appears in the upper pane and the graphical view appears in the lower pane.

2. In the graphical view, right-click the unallocated disk space on Disk 0 and choose New Partition. The New Partition Wizard appears.

3. Create a Primary Partition that is 250 MB. Accept the default drive letter assignment. Label the volume Data_Volume and perform a quick format using NTFS.

   After a few moments, a new drive named Data_Volume (*drive_letter*:) appears, where *drive_letter* is the letter that the New Partition Wizard assigned to the partition. When formatting is complete, the status of the partition displays Healthy.

Exercise 2: Converting a Disk from Basic to Dynamic Storage from the Disk Management Snap-In

1. In Disk Management, right-click Disk 0's status box in the graphical view and click Convert To Dynamic Disk. The Convert To Dynamic Disk dialog box appears and the Disk 0 check box is selected.
2. Follow the prompts to convert Disk 0 to a dynamic disk. Because Disk 0 is your system drive, the computer requires a restart.

Exercise 3: Using DiskPart

1. Open a command prompt.
2. Type **diskpart** and press Enter. The DISKPART> prompts appears.
3. Type **?** and press Enter. A list of Diskpart commands appear.
4. Type **list disk** and press Enter. A list of the disk or disks in Server01 appears.
5. Type **create volume simple size = 250 disk = 0** and press Enter.
6. Type **list volume** and press Enter.

   A new volume has been created. The new volume appears with an asterisk before its name. The asterisk denotes that the volume is selected. Notice that there is no drive letter assigned to the volume.
7. Type **assign letter z** and press Enter.
8. Type **list volume** and press Enter. The letter Z is assigned to the selected volume.
9. Type **extend size=250 disk=0** and press Enter.
10. Type **list volume** and press Enter. The selected volume (drive Z) is now 500 MB in size.
11. Type **exit** and press Enter. The command prompt reappears.
12. Type **format z: /fs:NTFS /v:Extended_Volume /q** and press Enter. A warning message appears stating that all data will be lost on drive Z.
13. Press Y and then press Enter. A quick format with NTFS is performed on drive Z.
14. Type **exit** to close the command window.

Exercise 4: Extending Volumes Using Disk Management

1. Open Disk Management.
2. Right-click Extended_Volume and choose Delete Volume.
3. Confirm the deletion of the volume by clicking Yes.
4. Right-click Data_Volume and choose Extend Volume. The Extend Volume Wizard appears.
5. Click Next.
6. Change the amount of space being used to extend the volume to 500 MB.
7. Click Next.
8. Read the summary information. Click Finish.

Exercise 5: Drive Letters and Mounted Volumes

1. Right-click Data_Volume and choose Change Drive Letter And Paths.
2. Change the drive letter to X.
3. Right-click Data_Volume (X:) and choose Open. Windows Explorer opens.
4. Create a folder called Docs.
5. Close Windows Explorer.
6. Right-click unallocated space on Disk 0 and choose New Volume.
7. Create a simple volume using all remaining space on the disk. Instead of assigning a drive letter, mount the volume in the path X:\Docs. Format the volume NTFS and label the volume More_Space.
8. Open Windows Explorer and make sure that Status Bar is selected in the View menu. Examine the X: volume. How much free space is shown? What free space is reported when you open the Docs folder?

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. This question continues the scenario that was presented in question 1 of the review in Lesson 1. You have installed a new 200 GB disk drive. You configured it as a basic disk and created three primary partitions of 30 GB each to host the operating system, user home directories, and shared data. You configured an extended partition and two logical drives of 30 GB each to host applications installed on the machine and a software distribution point. There remains 50 GB of unallocated space on the disk. Several months later, you notice that three of the volumes are nearing capacity. You want to prepare for the likely event that one or more partitions will need to be expanded. What action must you take?

   _____

   _____

2. What type of disk region supports logical drives?

   a. Primary partitions

   b. Simple volumes

   c. Spanned volumes

   d. Extended partitions

   e. Unallocated space

3. You recently added a disk to a computer. The disk had previously been used in a Windows 2000 Server. The disk appears in Device Manager, but is not appearing correctly in Disk Management. What task must you apply?

   a. Import Foreign Disk

   b. Format volume

   c. Rescan

   d. Change Drive Letter or Path

   e. Convert to Dynamic Disk

4. You attempt to convert an external FireWire disk from basic to dynamic, but the option to convert is not available. What is the most likely reason for this?

   _____

   _____

Lesson Summary

- Disk management tasks can be completed using the Disk Management snap-in or from the command prompt with tools such as Diskpart.
- Common disk management tasks include creating and deleting partitions and volumes and assigning drive letters and mount points.
- Windows Server 2003 allows you to assign one or no drive letter to a volume and, optionally, to mount the volume to one or more empty folders on NTFS volumes.
- Basic disks can be converted to dynamic disks, but all data and volumes must be deleted to convert a dynamic disk to a basic disk.

Lesson 3: Maintaining Disk Storage Volumes

Windows Server 2003 disk volumes are efficient and stable if formatted with NTFS, but somewhat less so when formatted with FAT or FAT32. The NTFS file system logs all file transactions, replaces bad clusters automatically, and stores copies of key information for all files on the NTFS volume. With these mechanisms, NTFS actively protects the integrity of the volume structure and the file system metadata (the data related to the file system itself). User data, however, can occasionally be corrupted, and can certainly become fragmented. Users also have the annoying habit of storing *enormous* amounts of archaic and non-business data on volumes to which they have access. In this lesson, you will learn how to maintain the integrity of disk volumes and to optimize those volumes by performing defragmentation and by setting storage limits through disk quotas.

After this lesson, you will be able to

- Monitor and maintain disk integrity using CHKDSK
- Monitor and improve disk performance using Disk Defragmenter
- Configure and monitor user disk storage using Disk Quotas

Estimated lesson time: 20 minutes

CHKDSK

CHKDSK, or "Check Disk", is a tool available in Windows Explorer or from the command-line that allows you to scan a disk volume for file system errors and, optionally, to test for and attempt to recover bad sectors on your hard disk.

To use Check Disk from Windows Explorer, open the properties dialog box for the volume you want to check. On the Tools tab, click Check Now. In the Check Disk dialog box, as shown in Figure 11-4, select the tasks you wish to launch.

Figure 11-4 The Check Disk dialog box

When you select Automatically Fix File System Errors, Check Disk will attempt to fix inconsistencies in the file system catalog, such as files that appear in the catalog but

don't appear in a directory on the volume. Check Disk makes three passes over the drive to examine the metadata, which is the data describing how files are organized on the disk. The passes attempt to ensure that all files on the volume are consistent with the master file table (MFT), that the directory structure is correct, and that the security descriptors are consistent.

If you select Scan For And Attempt Recovery Of Bad Sectors, Check Disk makes a fourth pass which tests the sectors in the volume reserved for user data (as opposed to file system metadata, which is always checked). If a bad sector is found, data is recovered and moved to a good sector if the volume is fault-tolerant; if the volume is not fault-tolerant, data cannot be recovered using Check Disk and must be restored from backup. The bad sector is then removed from active use and future data will not be written to the sector.

All files with open handles must be closed before Check Disk can run. If all handles cannot be released (which will be the case if you run Check Disk against a system volume), you will be prompted to schedule Check Disk to run when the system is restarted. When Check Disk is running, the volume will be inaccessible to other processes. Depending on the size of the volume, the check options you have selected, and the other processes running on the computer, Check Disk can take a significant amount of time to complete, and it is quite processor- and disk-intensive while it runs.

Check Disk can also be run from the command prompt using CHKDSK. Without switches, CHKDSK runs in read-only mode on the current drive. You'll see a report showing disk space usage. CHKDSK supports several switches allowing you to fix file system errors (/f) and bad sectors (/r), just like the Explorer version.

Disk Defragmenter

Files are stored on a volume in units called *clusters*. Cluster size is configured when formatting a drive; many NTFS volumes use a default cluster size of 4 KB. Each cluster can only contain one file, even if that file is smaller than the cluster size. If a file is larger than the cluster size, the file is saved to multiple clusters, with each cluster containing a pointer to the next segment of the file. When a drive is new, all clusters are free, so as files are written to the drive they tend to occupy physically adjacent clusters. But quickly, as files are deleted or expanded and contracted in size, free clusters are no longer completely contiguous, so a file may be saved to several clusters that are not physically close to each other on the disk drive. This fragmentation of a file results in slower read and write performance and, over time, fragmentation of multiple files on a server can degrade performance significantly.

Windows Server 2003 provides a defragmenter toolset—both a command-line and a graphical utility—with which volumes can be analyzed and defragmented. The tools are significantly improved over Windows 2000, as they can now defragment volumes

with cluster sizes greater than 4 KB, and can defragment the master file table. You can use the tools to defragment any local disk volume. But to schedule defragmentation, or to defragment a remote volume, you must look for a third-party tool such as Diskeeper from Executive Software.

To use the built-in Disk Defragmenter, as shown in Figure 11-5, open the properties of a disk volume and, from the Tools tab, click Defragment Now. Alternatively, open the Disk Defragmenter snap-in in the Computer Management console or a custom Microsoft Management Console. Select a volume and click Analyze. The tool will display a recommendation. If the tool indicates that the volume is dirty, there may be corruption and CHKDSK should be run before defragmenting.

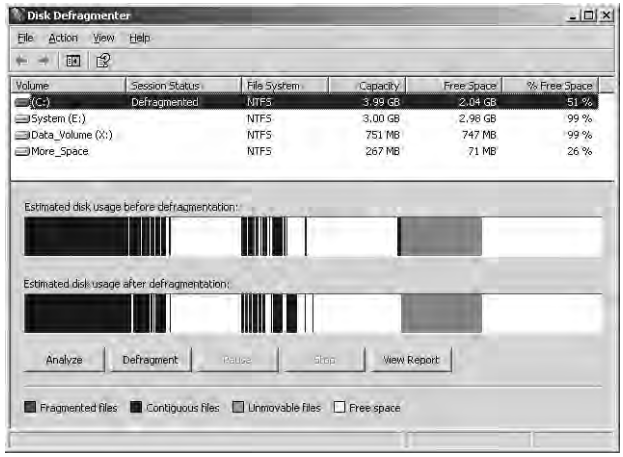

Figure 11-5 Disk Defragmenter

If the recommendation is to defragment, click Defragment. You can defragment any type of volume: FAT or NTFS, basic or dynamic. The volume can have open files, but open files may not be efficiently defragmented and may slow the process, so it is recommended to close all open files before defragmenting. Disk Defragmenter will move files around the drive in an attempt to collect all clusters of a file into contiguous clusters. The result will also consolidate free space, making it less likely that new files will be fragmented.

Note To completely defragment a volume, the volume must have at least 15 percent free space. This space is used to stage files as they are defragmented. If the volume contains numerous fragmented large files, the amount of free space required for effective defragmentation will be larger. If the volume contains less then 15 percent free space, then the volume will be only partially defragmented.

Disk Quotas

Windows 2000 introduced quota management as a built-in feature, allowing administrators to implement storage limits without an investment in third-party utilities. Windows Server 2003 supports the same functionality. When quotas are enabled, quota manager tracks the files on a volume that are owned by a user. It then compares the calculated total of disk usage by that user to limits that have been configured by an administrator and, when those limits are reached, notifies the user that the volume is near quota, or prevents the user from writing to the disk, or both.

Quota manager reports the amount of free space on a volume based on the user's quota, so if a user has a 50 MB quota on a 500 GB RAID volume, the user will see free space reported as 50 MB when the user first accesses the volume. When the user approaches the quota limit, the messages that appear are similar to a volume that is filling up or is full; the system warns that space is low and suggests deleting unneeded files.

Exam Tip Quotas are supported only on NTFS volumes.

Configure Quotas

Configuring quotas requires the following steps: enabling quotas on a volume, configuring default quota settings, and configuring quota entries for exceptions to the default.

Quotas are disabled by default in Windows Server 2003, and must be enabled on a volume-by-volume basis. To enable quotas, open the properties of the volume and click the Quota tab. The Quota properties of a volume are shown in Figure 11-6.

Tip Most documentation suggests opening the properties of the volume from Explorer, by right-clicking a drive and choosing Properties. Unfortunately, that process limits you to configuring quotas for lettered volumes only; Explorer cannot display the Quota tab for a volume mounted to a folder path. Therefore, it is recommended that you configure quotas from Disk Management. The Disk Management tool allows you to open the properties of any volume and access its Quota tab.

Select the Enable Quota Management check box. If you want to deny users who have exceeded their limit the ability to write additional files to the volume, select Deny Disk Space To Users Exceeding Quota Limit. If this box is not selected, users can continue to write to the volume.

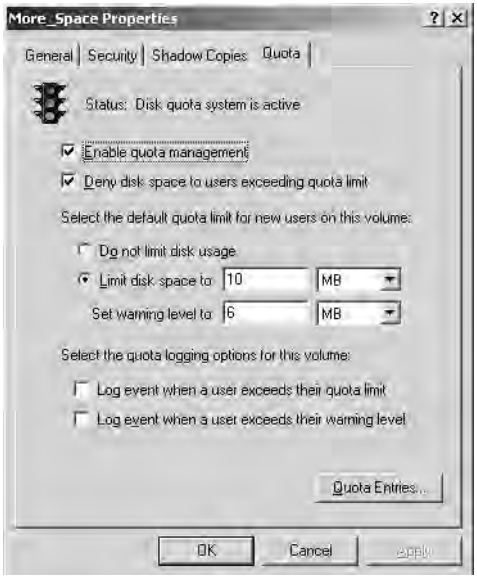

Figure 11-6 The Quota tab of a volume's Properties dialog box

Quotas are managed in two ways: first, by quota entries for specific users, setting a storage limit for each user (or setting "no limit" for a user), and second, by default quota settings that apply to all users for whom a quota entry does not exist. On the Quota tab, you can configure the default quota settings. Configure a default limit or "no limit" that will apply to as many users as possible, so that you can minimize the number of quota entries you must create for users whose limits are different from this default. Note that you can configure the disk space limit as well as a warning level, which should obviously be lower than the limit.

Finally, specify logging options. Quota manager registers events in the System log, identifying the user by name and specifying that they have exceeded their warning or quota limits.

After configuring the defaults for the volume on the Quota tab, click Quota Entries to open the Quota Entries dialog box, as shown in Figure 11-7.

Exam Tip Administrators have No Limit configured as their quota entry. That enables administrators to install the operating system, services, applications, and data without exceeding a quota.

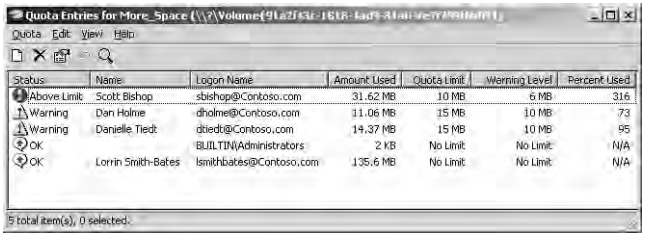

Figure 11-7 The Quota Entries dialog box

Click the New, Quota Entry button on the toolbar or choose New Quota Entry on the Quota menu, and you can select one or more users for which to create a quota entry. It is unfortunate that Windows Server 2003 does not allow you to assign quota entries based on groups (as most third-party quota management tools do), but in the Select Users dialog box you can at least select multiple users before clicking OK. The limits you configure in the Add New Quota Entry dialog box will apply to all selected users, individually.

Exporting Quota Entries

If you want to apply the same quota entries to another NTFS volume, you can export the entries and import them to the other volume. Select one or more quota entries and, on the Quota menu, click Export. On the other volume, choose Import.

Monitoring Quotas and Storage

The Quota Entries dialog box displays disk storage per user and whether that storage is at or above warning levels or limits. You can sort by column to identify users who have exceeded their quota levels or limits. There is no mechanism to alert you about quota limits, so you must monitor the Quota Entries dialog box or the System Log in Event Viewer.

Exam Tip Disk Quotas can be implemented per-volume and per-user only. You cannot implement quotas on a per-folder or per-group basis.

Practice: Implementing Disk Quotas

In this practice, you will configure default quota management settings to limit the amount of data users can store Server01. You will then configure custom quota settings to allow the users in the Marketing department to store more data, because their media files are generally larger than other users' business documents. And you allow developers to be exempt from quotas.

Exercise 1: Configuring Default Disk Quota Settings

1. Open Disk Management.
2. Right-click the More_Space volume and choose Properties.
3. Click the Quota tab.
4. Click the Enable Quota Management check box.
5. Select the Deny Disk Space To Users Exceeding Quota Limit check box.
6. Select Limit Disk Space To.

   Configure the limit as 10 MB and the warning level as 6 MB.
7. Select both Log check boxes.
8. Click Apply.

   A Disk Quota dialog box appears, warning you that the volume will be rescanned to update disk usage statistics if you enable quotas. Click OK to confirm.
9. Do not close the Volume Properties dialog box because you will use it in the next exercise.

Exercise 2: Creating Custom Quota Entries for Users

1. On the Quota tab of the More_Space Properties dialog box, click Quota Entries to open the Quota Entries dialog box.

Off the Record Notice that the Builtin\Administrators group is listed. If you created files while logged in as a non-administrator user, there would be a quota entry for that user as well because the user account owns files on the volume.

   You will now create quota entries allowing your Marketing employees, Dan Holme and Danielle Tiedt, more disk storage than the default.
2. On the Quota menu, click New Quota Entry.
3. Click Advanced and click Find Now. All users in the domain are listed.
4. Select Dan Holme and Danielle Tiedt and click OK twice.
5. Set the quota entry to limit disk usage at 15 MB, and to warn the user at 10 MB. Click OK.

   You will now create quota entries allowing your developers, Lorrin Smith-Bates and Scott Bishop, to be exempt from quotas.
6. Repeat steps 2 through 5 to configure quota entries for Lorrin Smith-Bates and Scott Bishop. Set the entry so that their disk usage is not limited.

Exercise 3 (Optional): Test Disk Quotas

1. Log on as Danielle Tiedt.
2. Create a folder in the X:\Docs folder called Dtiedt.
3. Copy the Support folder from the Windows Server 2003 CD-ROM to the X:\Docs\Dtiedt folder. The Support folder is 11 MB, and is lower than Danielle Tiedt's quota. The copy completes successfully.
4. Log on as Dan Holme.
5. Create a folder in X:\Docs called Dholme.
6. Copy the Support folder from the Windows Server 2003 CD-ROM to the X:\Docs\Dholme folder. The folder is smaller than Dan Holme's quota limit, and completes successfully.
7. Copy the Valueadd folder from the Windows Server 2003 CD-ROM to the X:\Docs\Dholme folder. The folder is 6 MB, and therefore puts Dan Holme over his quota limit. The copy will be interrupted.
8. Log on as Administrator and open the Quota Entries dialog box for the More_Space volume. Notice the information presented about disk usage for each user.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're the administrator of a Windows Server 2003 computer. You want to fix any file system errors and recover any bad sectors on your computer's hard disk. Which tool should you use?
   a. Check Disk
   b. Disk Defragmenter
   c. DISKPART
   d. Disk quotas

2. You're the administrator of a Windows Server 2003 computer. The computer's hard disk contains two data volumes: D and E. You enable disk quotas on volume D and E that limit all users to 20 MB of total storage. However, you want to limit storage in the users' home folders, stored in D:\Users, to 10 MB per user. Is this possible? Why or why not? Where can you implement quotas?
   a. On any server for all disks
   b. On any physical disk for all volumes
   c. On any volume for all folders
   d. On any folder
3. What is the required amount of free disk space on a volume in order to provide for complete defragmentation?
   a. 5 percent
   b. 10 percent
   c. 15 percent
   d. 25 percent
   e. 50 percent

Lesson Summary

- The Check Disk tool allows you to fix file systems errors and scan for and attempt to recover bad sectors on your hard disk.
- Disk Defragmenter improves performance by relocating files so that their clusters are contiguous.
- Disk Quotas allow you to set and monitor storage limits and, optionally, to deny write access to users exceeding those limits. Quotas are configured on a per-user, per-volume basis.

Lesson 4: Implementing RAID

A disk subsystem that includes a RAID configuration enables the disks in the system to work in concert to improve performance, fault tolerance, or both. In this lesson, you will learn about the three levels of RAID that can be created and managed by Windows Server 2003. You will learn the impact that each type of volume has on performance, volume capacity, and fault tolerance, and how to recover data in the event of a disk failure in a RAID configuration.

After this lesson, you will be able to

- Identify the best RAID implementation given a particular storage requirement regarding capacity utilization, fault tolerance, and performance
- Configure a striped volume (RAID-0)
- Configure a mirrored volume (RAID-1)
- Configure a RAID-5 volume (striped with parity)
- Recover from a single-disk failure in a fault-tolerant volume

Estimated lesson time: 25 minutes

Lesson 1 introduced the types of storage units available on a Windows Server 2003 computer. The types of volumes that reflect RAID configurations are striped volumes, mirrored volumes, and RAID-5 volumes.

Implementing Disk Fault Tolerance

As mentioned in Lesson 1, fault tolerance is the ability of a computer or operating system to respond to a catastrophic event, such as a power outage or hardware failure, so that no data is lost and that work in progress is not corrupted. Fully fault-tolerant systems using fault-tolerant disk arrays prevent the loss of data. You can implement RAID fault tolerance as either a hardware or software solution.

Hardware Implementations of RAID

In a hardware solution, the disk controller interface handles the creation and regeneration of redundant information. Some hardware vendors implement RAID data protection directly in their hardware, as with disk array controller cards. Because these methods are vendor specific and bypass the fault tolerance software drivers of the operating system, they offer performance improvements over software implementations of RAID.

Consider the following points when deciding whether to use a software or hardware implementation of RAID:

- Hardware fault tolerance is more expensive than software fault tolerance and might limit equipment options to a single vendor.
- Hardware fault tolerance generally provides faster disk I/O than software fault tolerance.
- Hardware fault tolerance solutions might implement hot swapping of hard disks to allow for replacement of a failed hard disk without shutting down the computer and hot sparing so that a failed disk is automatically replaced by an online spare.

Software Implementations of RAID

Windows Server 2003 supports one RAID implementation (striped, RAID-0) that is not fault-tolerant and two implementations that provide fault tolerance: mirrored volumes (RAID-1) and striped volumes with parity (RAID-5). You can create fault-tolerant RAID volumes only on dynamic disks formatted with NTFS.

With Windows Server 2003 implementations of RAID, there is no fault tolerance following a failure until the fault is repaired. If a second fault occurs before the data lost from the first fault is regenerated, you can recover the data only by restoring it from a backup.

Striped Volumes

A striped volume, which implements RAID Level 0, uses two or more disks and writes data to all disks at the same rate. By doing so, I/O requests are handled by multiple spindles, and read/write performance is the beneficiary. Striped volumes are popular for configurations in which performance and large storage area are critical, such as computer-aided design (CAD) and digital media applications.

Note You might not experience a performance improvement on IDE unless you use separate controllers. Separate controllers—ideally, one for each drive—will improve performance by distributing I/O requests among controllers as well as among drives.

Creating a Striped Volume

To create a striped volume, you must have unallocated space on at least two dynamic disks. Right-click one of the spaces and choose Create Volume. The New Volume Wizard will step you through the process of selecting a striped volume and choosing other disk space to include in the volume. Striped volumes can be assigned a drive letter and folder paths. They can be formatted only with NTFS.

Up to 32 disks can participate in a striped volume. The amount of space used on each disk in the volume will be equal to the smallest amount of space on any one disk. For example, if Disk 1 has 200 GB of unallocated space, and Disk 2 has 120 GB of space, the striped volume can contain, at most, 240 GB as the size of the stripe on Disk 1 can be no greater than the size of the stripe on Disk 2. All disk space in the volume is used for data; there is no space used for fault tolerance.

Recovering a Striped Volume

Because data is striped over more than one physical disk, performance is enhanced, but fault tolerance is decreased—there is more risk because if any one drive in the volume fails, all data on the volume is lost. It is important to have a backup of striped data. If one or more disks in a striped volume fails, you must delete the volume, replace the failed disk(s) and recreate the volume. Then you must restore data from the backup.

Exam Tip Striped volumes provide maximum storage and performance but support no fault tolerance. The only recovery potion is that of your regular backup routine.

Mirrored Volumes

A mirrored volume provides good performance along with excellent fault tolerance. Two disks participate in a mirrored volume, and all data is written to both volumes. As with all RAID configurations, use separate controllers (by adding a controller, you create a configuration called "duplexing") for maximum performance. Mirrored volumes relate to RAID-1 hardware configurations.

Create Mirrored Volumes

To create a mirrored volume, you must have unallocated space on two dynamic disks. Right-click one of the spaces and choose Create Volume. The New Volume Wizard will step you through the process of selecting a mirrored volume and choosing space on another disk to include in the volume. Mirrored volumes can be assigned a drive letter and folder paths. Both copies of the mirror share the same assignment.

You can also mirror an existing simple volume by right-clicking the volume and choosing Add Mirror and selecting a drive with sufficient unallocated space.

Once you have established the mirror, the system begins copying data, sector by sector. During that time, the volume status is reported as Resynching.

Recovering from Mirrored Disk Failures

The recovery process for a failed disk within a mirrored volume depends on the type of failure that occurs. If a disk has experienced transient I/O errors, both portions of

the mirror will show a status of Failed Redundancy. The disk with the errors will report a status of Offline or Missing, as seen in Figure 11-8.

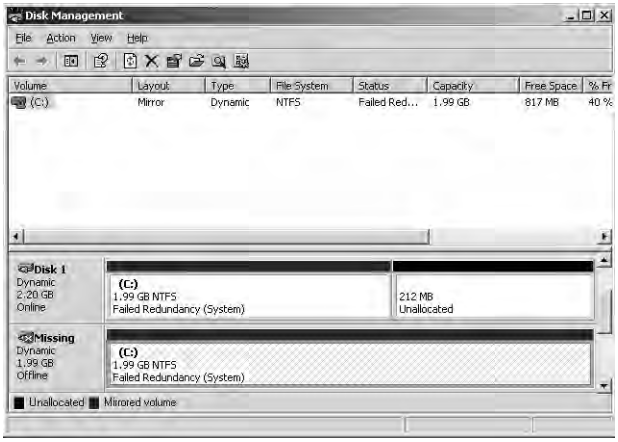

Figure 11-8 A mirrored volume with a failed disk

After correcting the cause of the I/O error—perhaps a bad cable connection or power supply—right-click the volume on the problematic disk and choose Reactivate Volume or right-click the disk and choose Reactivate Disk. Reactivating brings the disk or volume back online. The mirror will then resynchronize automatically.

If you want to stop mirroring, you have three choices, depending on what you want the outcome to be:

- **Delete the volume** If you delete the volume, the volume and all the information it contains is removed. The resulting unallocated space is then available for new volumes.
- **Remove the mirror** If you remove the mirror, the mirror is broken and the space on one of the disks becomes unallocated. The other disk maintains a copy of the data that had been mirrored, but that data is of course no longer fault-tolerant.
- **Break the mirror** If you break the mirror, the mirror is broken but both disks maintain copies of the data. The portion of the mirror that you select when you choose Break Mirror maintains the original mirrored volume's drive letter, shared folders, paging file, and reparse points. The secondary drive is given the next available drive letter.

Knowing that information, how do you suppose you would replace a failed disk—a member of the mirrored volume that simply died? Well, after physically replacing the disk, you will need to open Disk Management to rescan, initialize the disk and convert it to dynamic. After all that work you will find that you can't remirror a mirrored volume, even though half of it doesn't exist. So far as the remaining disk is concerned, the mirrored volume still exists—its partner in redundancy is just out to lunch. You must

remove the mirror to break the mirror. Right-click the mirror and choose Remove Mirror. In the Remove Mirror dialog box, it is important to select the half of the volume that is missing; the volume you select will be deleted when you click Remove Mirror. The volume you did not select will become a simple volume. Once the operation is complete, right-click the healthy, simple volume and choose Add Mirror. Select the new disk and the mirror will be created again.

Exam Tip Mirrored volumes provide fault tolerance and better write performance than RAID-5 volumes. However, because each disk in the mirror contains a full copy of the data in the volume, it is the least efficient type of volume in terms of disk utilization.

RAID-5 Volumes

A RAID-5 volume uses three or more physical disks to provide fault tolerance and excellent read performance while reducing the cost of fault tolerance in terms of disk capacity. Data is written to all but one disk in a RAID-5. That volume receives a chunk of data, called parity, which acts as a checksum and provides fault tolerance for the stripe. The calculation of parity during a write operation means that RAID-5 is quite intensive on the server's processor for a volume that is not read-only. RAID-5 provides improved read performance, however, as data is retrieved from multiple spindles simultaneously.

As data in a file is written to the volume, the parity is distributed among each disk in the set. But from a storage capacity perspective, the amount of space used for fault tolerance is the equivalent of the space used by one disk in the volume.

From a storage capacity perspective, that makes RAID-5 more economical than mirroring. In a minimal, three disk RAID-5 volume, one-third of the capacity is used for parity, as opposed to one-half of a mirrored volume being used for fault tolerance. Because as many as 32 disks can participate in a RAID-5 volume, you can theoretically configure a fault-tolerant volume which uses only 1/32 of its capacity to provide fault tolerance for the entire volume.

Configure RAID-5 Volumes

You need to have space on at least three dynamic disks to be able to create a RAID-5 volume. Right-click one disk's unallocated space and choose New Volume. The New Volume Wizard will step you through selecting a RAID-5 volume type, and then selecting the disks that will participate in the volume.

The capacity of the volume is limited to the smallest section of unallocated space on any one of the volume's disks. If Disk 2 has 50 GB of unallocated space, but Disks 3 and 4 have 100 GB of unallocated space, the stripe can only use 50 GB of space on Disks 3 and 4—the space used on each disk in the volume is identical. The capacity,

or Volume Size reported by the New Volume Wizard will represent the amount of space available for data after accounting for parity. To continue our example, the RAID-5 volume size would be 100 GB—the total capacity minus the equivalent of one disk's space for parity.

RAID-5 volumes can be assigned a drive letter or folder paths. They can be formatted only with NTFS.

Because RAID-5 volumes are created as native dynamic volumes from unallocated space, you cannot turn any other type of volume into a RAID-5 volume without backing up that volume's data and restoring into the new RAID-5 volume.

Recovering a Failed RAID-5 Volume

If a single disk fails in a RAID-5 volume, data can continue to be accessed. During read operations, any missing data is regenerated on the fly through a calculation involving remaining data and parity information. Performance will be degraded and, of course, if a second drive fails it's time to pull out the backup tapes. RAID-5 and mirrored volumes can only sustain a single drive failure.

If the drive is returned to service, you may need to rescan, and then you will need to right-click the volume and choose Reactivate Volume. The system will then rebuild missing data and the volume will be fully functional again.

If the drive does not offer a Reactivate option, or if you have had to replace the disk, you may need to rescan, initialize the disk, convert it to dynamic, then right-click the volume and choose Repair Volume. You will be asked to select the disk where the missing volume member should be recreated. Select the new disk and the system will regenerate the missing data.

Mirrored Volumes versus RAID-5 Volumes

Mirrored volumes (RAID-1) and RAID-5 volumes provide different levels of fault tolerance. Deciding which option to implement depends on the level of protection you require and the cost of hardware. The major differences between mirrored volumes and RAID-5 volumes are performance and cost. Table 11-2 describes some differences between software-level RAID-1 and RAID-5.

Table 11-2 RAID Performance and Costs

Mirrored Volumes (RAID-1)	Striped Volumes with Parity (RAID-5)
Can protect system or boot partition	Cannot protect system or boot partition
Requires two hard disks	Requires a minimum of three hard disks and allows a maximum of 32 hard disks

Table 11-2 RAID Performance and Costs (Continued)

Mirrored Volumes (RAID-1)	Striped Volumes with Parity (RAID-5)
Has a higher cost per MB	Has a lower cost per MB
50 percent redundancy*	33 percent maximum redundancy*
Has good read and write performance	Has excellent read and moderate write performance
Uses less system memory	Requires more system memory

* drive space dedicated or "lost" to provide fault tolerance

Creating Fault Tolerance for the System Volume

Because RAID-5 is a native dynamic volume, it is not possible to install or start the Windows Server 2003 operating system on a RAID-5 volume created by the Windows Server 2003 fault-tolerant disk technologies.

Tip *Hardware RAID*, however, is invisible to Windows Server 2003, so the operating system can (and should, where available) be installed on hardware RAID arrays.

The only option for creating fault tolerance for the system, without buying hardware RAID, is thus to mirror the system volume. You can mirror the system volume by following the procedures described for creating a mirrored volume: right-click the system volume and choose Add Mirror. Unlike Windows 2000, you do not need to restart, and the BOOT.INI file is updated automatically so that you can start to the secondary drive if the primary drive fails.

If the drives are attached to IDE controllers, and the primary drive fails, you may have to remove that drive, change the secondary drive to the primary controller and set its jumpers or cable position so that it is the master. Otherwise, the system may not boot to the secondary drive.

Tip If you are going to mirror the system volume, do so on one or two SCSI controllers. If you use two controllers, make sure they are of the same type. This configuration will be the most easily supported and recovered.

Upgrading Disks

There are two potential "gotchas" when you upgrade disks from previous versions of Windows, or attempt to move disks to a Windows Server 2003 computer from a computer running a previous version of Windows.

First, if a disk was configured in a Windows 2000 computer as a basic disk, then was converted to dynamic, you cannot extend that disk's simple volumes onto other disks using Windows Server 2003. In other words, if you move that disk to a Windows Server 2003 computer, or upgrade the operating system to Windows Server 2003, you cannot create spanned volumes out of the disk's simple volumes.

Second, Windows Server 2003 no longer supports multidisk arrays created in Windows NT 4. Windows NT 4 created mirrored, striped, and striped-with-parity (RAID-5) sets using basic disks. Windows 2000 permitted the use of those disk sets, although it was important to convert the sets to dynamic quickly in order to facilitate troubleshooting and recovery. Windows Server 2003 does not recognize the volumes. On the off chance that you upgrade a server from Windows NT 4 to Windows Server 2003, any RAID sets will no longer be visible. You must first back up all data prior to upgrading or moving those disks, and then, after recreating the fault-tolerant sets in Windows Server 2003, restore the data.

Practice: Planning RAID Configuration

In this practice, you will evaluate a server and its storage capacity against the requirements of *contoso.com* and determine an appropriate configuration.

You administer a server for Contoso, Ltd. The server has four disks on a SCSI subsystem:

- Disk 0: 80 GB
- Disk 1: 80 GB
- Disk 2: 40 GB
- Disk 3: 40 GB

You recently performed a clean installation of Windows Server 2003 by backing up all data on the disks, removing all partitions from those disks, and installing the operating system on a 20 GB partition on Disk 0.

You are now required to configure all the remaining drive space. User data will not be stored on the operating system volume. You want to maximize data storage and ensure uptime in the event of a single disk failure. What configuration do you implement, and what will the total storage capacity for user data be?

The answer is a combination of RAID-5 and mirrored volumes with a total capacity for user data of 140 GB.

To ensure uptime in the event of a single disk failure, you must provide fault tolerance for the operating system itself. Only a mirrored volume is capable of doing that; you cannot install or host the operating system on a RAID-5 volume. A minimum disk space of 20 GB is therefore required to mirror the operating system.

A RAID-5 configuration maximizes disk space without sacrificing single disk failure fault tolerance. You can configure a RAID-5 volume with three or more disks. In this scenario, configuring a RAID-5 volume with all four disks would maximize data storage. A RAID-5 volume's stripe can only be as wide as the smallest amount of unallocated space, so although disk 0 and 1 have 60 and 80 GB free, respectively, the smaller (40 GB) drives will determine the capacity of the volume. With a 40 GB space on four drives, the volume has a potential capacity of 160 GB, but RAID-5 uses the space equivalent to one disk for parity, meaning that the resulting capacity for data storage in this volume will be 120 GB.

That leaves disk 0 with 20 GB of unallocated space, and disk 1 with 40 GB of unallocated space. You can configure the mirror of the operating system volume on disk 1, leaving 20 GB on that drive. The remaining space (20 GB per disk on disks 0 and 1) can be configured as a mirrored volume for user data, with a storage capacity of 20 GB. A simple, spanned, or striped volume would not be fault-tolerant, and a RAID-5 volume requires a minimum of three physical disks, so a mirror is the most effective way to use remaining space for fault-tolerant data storage.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're implementing software RAID on your Windows Server 2003 computer. You want to provide fault tolerance to the system and boot partitions. Which version of RAID should you use?

   a. RAID-0

   b. RAID-1

   c. RAID-5

   d. You cannot use software RAID to protect a boot partition.

2. You're setting up a Windows Server 2003 computer and you want to protect the data on the hard disk. You want to implement a solution that provides the fastest disk I/O possible and supports the hot swapping of hard disks. Which RAID solution should you use?

   a. RAID-0

   b. RAID-1

   c. RAID-5

   d. Hardware RAID

3. You're setting up RAID-5 on your Windows Server 2003 computer. You plan to use five hard disks, which are each 20 GB in size. What percentage of redundancy can you anticipate with this configuration?

   a. 20

   b. 25

   c. 33

   d. 50

4. You're setting up software RAID on your Windows Server 2003 computer to provide fault tolerance to the data stored on that system. The computer is used as a database server. The server performs many read operations but relatively few write operations. As a result, you want a fault-tolerant solution that provides excellent read performance. Which RAID solution should you use?

   a. RAID-0

   b. RAID-1

   c. RAID-5

5. A computer where you want to implement RAID-5 contains three disks, each with 2 GB of unallocated space. Using the Disk Management snap-in, you start the New Volume Wizard by right-clicking one of the regions of unallocated space. When you reach the Select Volume Type screen, the RAID-5 option is not available. What is the most likely reason for this behavior?

   a. RAID-5 is already implemented in hardware.

   b. One or two of the disks are configured with the basic storage type.

   c. All three disks are configured with the dynamic storage type.

   d. All three disks are configured with the basic storage type.

   e. RAID-5 is already implemented in software.

6. A disk in a mirrored volume is failing. You decide to replace the failing disk. How should you prepare the mirror for disk replacement?

Lesson Summary

- Some levels of RAID provide fault tolerance by implementing data redundancy. You can implement RAID fault tolerance as either a software or hardware solution.
- Hardware solutions offer better performance than software solutions, but they are generally more expensive.
- Windows Server 2003 supports three software implementations of RAID: striped volumes (RAID-0), mirrored volumes (RAID-1), and striped-with-parity volumes (RAID-5).
- A striped volume (RAID-0) distributes data across each disk in the volume, providing increased read and write performance, but no benefit to fault tolerance.
- In a RAID-5 volume, fault tolerance is achieved by adding a parity-information stripe to each disk partition in the volume.
- A mirrored volume uses the fault tolerance driver to write the same data to a volume on each of two physical disks simultaneously.
- The major differences between mirrored volumes and RAID-5 volumes are performance and cost. Mirrored volumes offer good read and write performance. RAID-5 volumes offer better read performance than mirrored volumes, but only moderate write performance.
- The only form of software RAID that can be used for the system volume is a mirrored volume.

Case Scenario Exercise

Note This case scenario requires Internet access.

You are a server administrator for Contoso, Ltd. The company's file servers are running out of disk capacity, and it is necessary to upgrade. In the past, the company has relied on tape backups for data redundancy. Due to recent growth, it is no longer acceptable to encounter more than a few minutes of downtime if a server disk drive fails. You have therefore been asked to evaluate disk storage options that provide fault tolerance.

Exercise 1: Consider Windows Server 2003 Fault-Tolerant Volumes

Review the information in Lesson 4 to consider how you could best configure fault-tolerant servers using Windows Server 2003 dynamic volumes. Use the Practice in Lesson 4 as a reminder of how various types of volumes can be configured to support fault tolerance.

Consider the challenges related to IDE drives. If the operating system is installed on a mirrored IDE drive and the primary drive fails, you must reconfigure the secondary drive's jumpers or cable position, and ensure it is attached to the primary IDE channel. With that in mind, you decide that a more robust configuration would utilize two SCSI controllers with one copy of the mirror as the first disk on each SCSI chain. That configuration would enable rapid recovery not only from a single drive failure, but from the failure of one of the SCSI controllers as well.

Now consider the performance and capacity effect of Windows Server 2003 RAID. Consider the amount of time that will be required to recover if a drive fails—downing the server, replacing the drive, restarting the server—and the amount of time it will take to regenerate a missing volume.

Exercise 2: Consider Hardware RAID

With all those thoughts in mind, you decide to examine hardware RAID as an option. What advantages does hardware RAID provide? See Lesson 4 for some of the answers.

Open Internet Explorer and browse to the Web site(s) of one or more computer hardware and supply vendors. Search their sites for RAID arrays. You will find RAID arrays, which include disk drives, and RAID controllers and RAID enclosures, to which you must add drives. Focus on the ready-to-go RAID arrays and answer the following questions:

- What options are available?
- What are some of the vendors of hardware RAID arrays?
- What types of storage capacities do hardware RAID arrays offer?
- What RAID configurations do the hardware RAID arrays implement? Are there configurations that Windows Server 2003 does not support?
- What is the price range for a hardware RAID array?
- What do some entry-level RAID arrays cost?

At the time of this writing, hardware RAID solutions offering 720 GB of storage—that's closer to a terabyte than to the size of any single drive in most servers—can be purchased for less than $3,000.

How would you position the value of hardware RAID to your manager? Would you recommend hardware RAID over Windows Server 2003 RAID? Why or why not?

Troubleshooting Lab

You are a server administrator for Contoso, Ltd. You inherited a server from a previous administrator that contains numerous internal SCSI disk drives. You open the Disk Management console to determine the configuration of those drives and their volumes. The configuration is shown below:

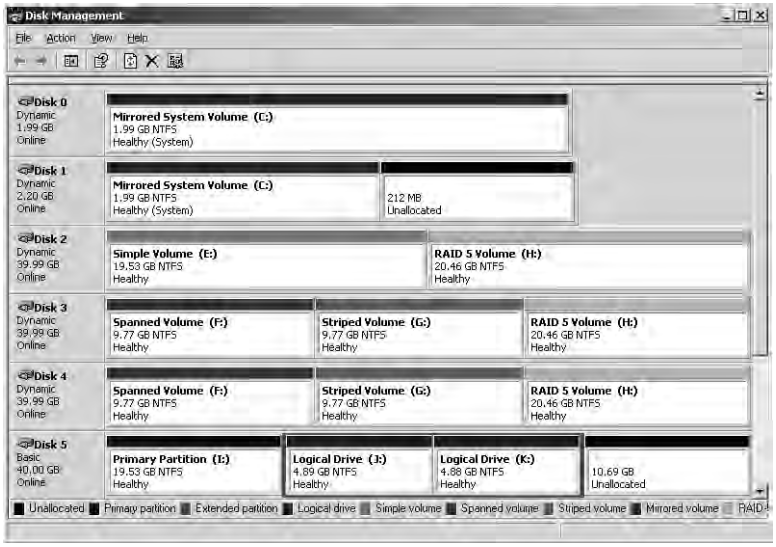

The weather forecast calls for a brutal storm to move into the city early tomorrow morning. To play it safe, you start a backup of your server on your way out the door. The storm is quite strong, forcing businesses, including yours, to be closed for several days. Electricity is lost to Contoso's building and, eventually, the batteries in your uninterruptible power supplies (UPSs) are drained, causing power to your servers to be completely lost. During the first few hours in which electricity is restored, several power fluctuations and surges are experienced.

When you return to the server room, you boot the servers. Your server indicates errors, and you open Disk Management to see the following, frightening graphical view of your disks and volumes:

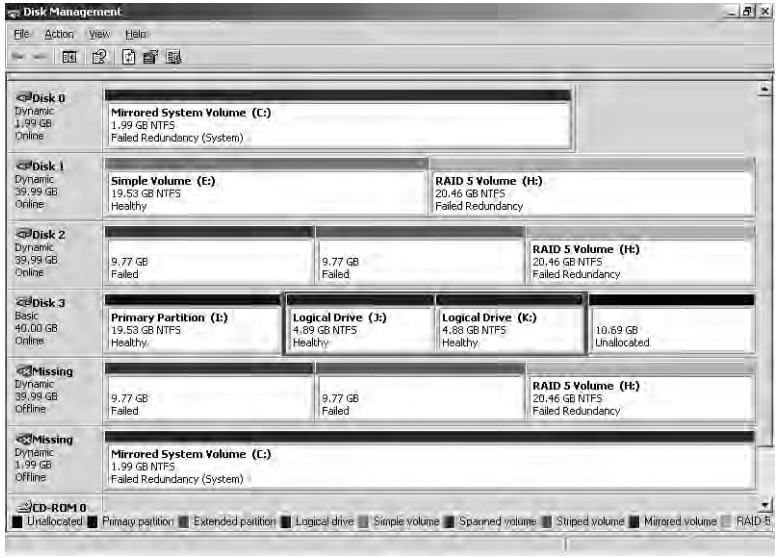

Two drives have failed in the server. One contained a mirror of the operating system volume. The other contained several volume types, including portions of a spanned, a striped, and a RAID-5 volume.

You have an 80 GB drive, still in its box. You shut down the server and remove the two failed drives. After inserting the new disk, you reboot the server.

Exercise

Take a moment, on a separate piece of paper, to plot the steps that will be required to recover the data on each volume that was lost. Be thorough. Include the steps required to clean up the missing disks and volumes as well as install, configure, and replace data on the new disk.

When you are confident that you have as comprehensive a list of steps as possible, compare your answer to the answer.

The components of recovery will include the following:

1. Log on to the system.
2. Finish installing the new disk drive. Follow any instructions presented by the Found New Hardware Wizard. If the Found New Hardware Wizard does not appear, check Device Manager to see if the disks installed automatically and silently. If the disks do not appear, use Add Hardware to install the disks.

3. Open Disk Management.
4. Detect and initialize the new disk. Disk Management will likely detect the new disk and present the Initialize Disk Wizard. If the wizard does not appear, check to see if the disk appears in Disk Management and, if not, right-click Disk Management and choose Rescan. Once the disk appears, right-click the disk and choose Initialize.
5. Recover the volumes (in any order).

   Recover the RAID-5 volume

   a. Convert the new disk to a dynamic disk. Right-click the new disk and choose Convert to Dynamic.
   b. Right-click a functioning portion of the RAID-5 volume and choose Repair Volume. Select the new disk, which has ample space to support a member of the stripe. The RAID-5 volume will be created and synchronized.

   Recover the mirrored volume

   a. Remove the mirror. Right-click the failed drive and choose Remove Mirror. Confirm that the portion marked Missing is selected and click Remove Mirror. The remaining portion of the mirror becomes a simple volume.
   b. Right-click the simple volume and choose Add Mirror. Select the new disk, which has ample space for the mirror, and click Add Mirror. The mirror will be created and synchronized.

   Recover the striped volume

   a. Delete the volume. Striped volumes are not fault-tolerant. All data on the volume was lost.
   b. Re-create the volume. Right-click on unallocated space where the stripe had existed, and choose New Volume. Select a striped volume and add the new disk to the stripe. The striped volume will be created and formatted.
   c. Restore data from the backup to the striped volume.

   Recover the spanned volume

   a. Delete the volume. Spanned volumes are not fault-tolerant. All data on the volume was lost.
   b. Re-create the volume. Right-click on unallocated space where the volume had existed, and choose New Volume. Select a spanned volume and add the new disk to the stripe. Select the appropriate amount of space to use on the new disk. The spanned volume will be created and formatted.
   c. Restore data from the backup to the spanned volume.

6. Remove the missing disks. Right-click the missing disks and choose Remove Volume. You cannot remove the disk with the missing mirror until after the mirror has been removed. You cannot remove the disk with the simple, spanned, and RAID-5 volumes until those volumes have been deleted and repaired.
7. Run CHKDSK after all volumes have been resynchronized and restored.

Chapter Summary

- Windows Server 2003 supports two types of storage, basic and dynamic, and several file systems, including FAT, FAT32, and NTFS. Most advanced storage management features are available only on dynamic disk volumes formatted as NTFS.
- Dynamic disks provide flexible and powerful options in configurations with more than one disk. You can implement spanned, mirrored, striped, and RAID-5 volumes to provide storage according to capacity, performance, and fault tolerance requirements.
- Disk volumes can be corrupted, can become fragmented, and often fill to capacity. Check Disk, Disk Defragmenter, and Disk Quotas are tools to help you manage existing volumes.
- Not all RAID configurations are fault-tolerant—mirrored and RAID-5 volumes are fault-tolerant, but striped volumes are not. None of the Windows Server 2003 volume types will provide fault tolerance if more than one disk fails in the volume.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Understand the impact on capacity, performance, and fault tolerance for each type of disk volume. Be prepared to recommend disk configurations based on storage requirements.
- Know how to implement user disk quotas and the effect of both default quota settings and specific quota entries.
- Recognize and repair a volume that was temporarily offline but is now reconnected: Rescan and Reactivate Disk or Reactivate Volume, and CHKDSK.
- Know how to rebuild fault-tolerant volumes (mirrored and RAID-5 volumes) on a replaced disk and the appropriate commands that are used: Rescan, Initialize, Convert to Dynamic Disk, Break Mirror, Remove Mirror, and Repair Volume.

Key Terms

Simple volume The equivalent to a basic disk partition is a dynamic disk simple volume. Because simple volumes exist on only one physical disk, they are not fault-tolerant.

Spanned volume A spanned volume includes space on more than one physical disk. Because their size tends to be greater, and because multiple physical disks are involved, the risk for failure increases, and spanned volumes are not fault-tolerant.

Striped volume Data is written to 2 to 32 physical disks at the same rate. Offers maximum performance and capacity but no fault tolerance.

Mirrored volume Two disks contain identical copies of data. The only software RAID supported on the system volume. Good read and write performance; excellent fault tolerance; but costly in terms of disk utilization, because 50 percent of the volume's potential capacity is used for data redundancy.

RAID-5 volume Data is written to 3 to 32 physical disks at the same rate, and is interlaced with parity to provide fault tolerance for a single disk failure. Good read performance; good utilization of disk capacity; expensive in terms of processor utilization and write performance as parity must be calculated during write operations.

12 Monitoring Microsoft Windows Server 2003

Exam Objectives in this Chapter:

- Monitor current system performance
- Monitor and analyze events
- Monitor and optimize a server for application performance
- Monitor memory performance objects
- Monitor network performance objects
- Monitor process performance objects
- Monitor disk performance objects
- Monitor server hardware for bottlenecks
- Monitor events

Why This Chapter Matters

When you first install a new computer, full of resources and uncomplicated by time and use, much is right with the world. However, as the newness of your server fades, and more demands are placed upon it by added applications and users, problems can develop. Without knowledge of the monitoring tools available and the best way to use them in your environment, you may watch your server performance degradation go from an annoyance to a significant problem.

The first monitoring steps for a new Microsoft Windows Server 2003 computer should include a thorough baselining of resource availability and performance data, which should be compared periodically with real-time data so that developing problems with applications or hardware can be solved or averted before they become serious. With the broad range of tools available in Windows Server 2003, no self-respecting system administrator should be caught unaware.

Lessons in this Chapter:

Before You Begin

To follow and perform the practices in this chapter, you need:

- A computer named Server01 with Windows Server 2003 installed.
- Server01 should be configured as a domain controller in the *contoso.com* domain.

Lesson 1: Using Event Viewer

Windows Server 2003 includes a set of log files that are configured and presented within the Event Viewer. By configuring the options on each of the logs to meet the requirements of your environment, you can collect data appropriate for troubleshooting hardware, application, system, and resource access.

After this lesson, you will be able to

- Identify the types of Event Viewer Logs
- Configure the appropriate recording of log data
- Display logged data in filtered form

Estimated lesson time: 20 minutes

Logs Available in Event Viewer

The Windows Server 2003 Event Log service, present and started automatically on all Windows Server 2003 computers, records events in one of three log files:

- **Application** Developers of an application can program their software to report configuration changes, errors, or other events to this log.
- **System** The Windows Server 2003 operating system will report events (service start or abnormal shutdown, device failures, and so on) to this log. The events reported to this log are preconfigured.
- **Security** Logon and resource access events (audits) are reported to this log. Configuration for most of these events is at the discrimination of the system administrator.

Note Although the Application and System log events are determined by the application developer and operating system, respectively, the Security log must first be configured for the type of events to record (Success or Failure for each). If File and Object Access events are selected, the security properties of each object must be configured to record auditing events to the Security log.

Windows Server 2003 computers filling the role of a Domain Controller contain two additional logs:

- **Directory Service** This log contains events related to the Microsoft Active Directory directory service, such as irreconcilable object replication or significant events within the directory.

- **File Replication Service** This log contains errors or significant events reported by the File Replication Service related to the copying of information between Domain Controllers during a replication cycle.

Lastly, a Windows Server 2003 computer filling the role of a Domain Name System (DNS) server will contain one additional log:

- **DNS Server** This log contains errors or significant events reported by the DNS server.

Configuring Event Viewer Logs

When you first start Event Viewer, all events that are recorded in the selected log are displayed. Such a list may be lengthy, containing many entries of both informational and warning types. You can locate events by type using the Filter command on the shortcut menu's View menu for the log you want to view. The Filter properties page for the Security log is shown in Figure 12-1.

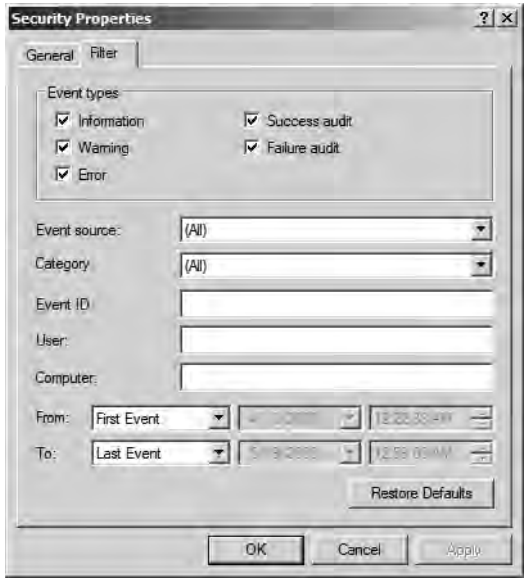

Figure 12-1 Filter settings for the Security log

Adjacent to the Filter tab in the properties of a log is the General tab, which provides access to the behaviors of the log, including

- The display name for the view of the log.
- The maximum size of the log.

- Whether the oldest events in the log should be overwritten when the maximum log size is reached. There are three overwrite options:.
  - **Overwrite Events As Needed (default)** This behavior will overwrite the oldest entries in the log with newer ones when the log reaches the maximum size.
  - **Overwrite Events Older Than *n* Days** This configuration will overwrite events that exceed the age setting when the log reaches the maximum size.
  - **Do Not Overwrite Events (Clear Log Manually)** This configuration will halt event logging when the log reaches the maximum size.

Security Alert Leaving the default setting of Overwrite Events As Needed on the Security log could overwrite important resource access or other security-related data if the log is not checked often. A regular schedule of analysis is recommended. Log files can be archived (that is, saved to disk) if needed for record-keeping or other administrative purposes.

For better assurance that no Security log entries have been lost, Windows Server 2003 Group Policy provides a setting in the Computer Configuration Policy: Security Settings that will force a computer to shutdown if it is unable to write to the Security log with audit information. This setting forces disciplined administrative practice if the Security log is set to be cleared manually.

The General tab for the Security log is shown in Figure 12-2.

Figure 12-2 The General settings for the Security log

Practice: Event Monitor

In this practice, you will configure the Security log for File and Object Access, and filter the data displayed in the Security log.

Exercise 1: Configuring the Security Log

In this exercise, you will configure the auditing of File and Object Access.

1. Logged on to Server01 as an administrator, open Active Directory Users And Computers.
2. Right-click the Domain Controllers Organizational Unit (OU), and then choose Properties from the shortcut menu.
3. On the Group Policy tab, select the Default Domain Controllers Policy, and then click Edit.
4. Under the Computer Configuration node, expand Windows Settings, Security Settings, Local Policies, and then click Audit Policy.
5. In the details pane, right-click Audit Object Access, and then select Properties from the shortcut menu.
6. In the Audit Object Access Properties dialog box, select Audit These Attempts: Failure, and then click OK.
7. Close the Group Policy Object Editor, click OK to close the Domain Controllers Properties dialog box, and then close Active Directory Users And Computers.
8. Open a command window, type **gpupdate,** and then press Enter.
9. When the Computer Policy reports as refreshed, close the command window.

You have now enabled the auditing of failed Object Access attempts on Server01 (as part of the Domain Controllers OU), and refreshed Group Policy so that the settings take effect immediately.

Exercise 2: Setting File and Object Auditing

In this exercise, you will configure auditing on a folder that you will create. Permissions will be set so as to simulate a user attempting to gain unauthorized access to the resource.

1. On your desktop, create a folder called Data.
2. Right-click the folder and select Properties from the shortcut menu.
3. Select the Security tab, and then select your user account.
4. Select the check box indicating Deny:Full Control permissions for your user account, click Yes in the warning dialog box.

5. Click Advanced, and then select the Auditing tab. Add your user account to audit List Folder / Read Data: Failed, and then click OK to close all Property dialog boxes.
6. Double-click the Data folder to open it. You should receive an Access Denied warning message.

Exercise 3: Reading the Security Log

In this exercise, you will confirm the auditing of your failed access to the Data folder.

1. From Administratives Tools, open the Computer Management console.
2. Expand the Event Viewer node, and then click the Security log in the folder pane.

   Near the top of the list of events, you should see several Failure Audit events (with ID 560) indicating your failed attempt to access the Data folder.
3. Right-click the Security log in the folder pane, select View from the shortcut menu, and then choose Filter.
4. In the Filter dialog box, select each of the following:
   - Event Source: Security
   - Category: Object Access
   - Event Types: Failure selected, all others cleared
5. Click OK to apply the filter to the Security log.

You have now filtered the Security log data to display only the events that apply to failed object access.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. On a domain controller running DNS, what logs will Event Viewer display by default? What are these logs, and what data do they collect?

   _____

   _____

2. You have configured your Windows Server 2003 computer to audit all failed object access, and all files and folders have auditing configured for List Folder / Read Data Failure. All other Event Viewer and Security log settings are at their default configurations. What will happen when the number of entries in the Security log reaches 512 KB?

3. You do not want data in the Security log to be overwritten, but also do not want your Windows Server 2003 computer to stop serving the network at any time. What settings will you configure on your server?

Lesson Summary

The Windows Server 2003 Event Viewer contains several logs which report errors and significant events during system operation. The System log contains data related to service and other internal operating system functioning. The Application log contains data written to it by software programs. The Security log contains data for successful and failed audits. Different from the other logs, the Security log is configurable by the administrator as to what data is written to it. Domain Controllers have additional logs for File Replication Services and the Active Directory. DNS servers have a log for DNS that is separate from other application logs.

In complex or lengthy logs, the display of data can be filtered by various criteria, including recording date and type of data, to make the data more readable. The data from the logs can be stored in a variety of file types, as follows:

- Performance Monitor Binary file (*.blg), with or without circular overwriting
- Text file (*.txt or *.csv)
- SQL Database

Once the settings for a log have been configured, the Performance Monitor settings can be saved (by selecting Save Settings As, from the log's context menu) as a hypertext markup language (HTML) file for later use (by selecting New Log Settings From, from the Counter Logs context menu).

All Event Viewer logs can be configured separately as to the maximum size of the log file, and how that log file should operate if the maximum size is reached. The choices when the maximum file size is reached is to overwrite older data immediately, only overwrite data that is of a certain age, or force a manual clearing of the log and never overwriting any data. In a related configuration, Group Policy can be enabled so as to force the immediate shutdown of a computer that is unable to write audit information to the Security log.

Lesson 2: Using the Performance Console

With the Performance Console, you can measure the activity of any computer on the network. System Monitor and the Performance Logs And Alerts snap-ins are built in as parts of the Performance console (perfmon.msc). The System Monitor snap-in allows for the viewing of real-time performance data as collected from configurable counters. The Performance Logs And Alerts snap-in allows for the recording of performance data (logs) and configurable actions when a threshold for a counter is breached (alerts). The Performance console allows you to perform multiple tasks, including the following:

- Collect and view real-time performance data
- View data collected in a log
- Present data in a graph, histogram, or report view
- Create HTML pages from views by importing of log file settings
- Save monitoring configurations that can be loaded into System Monitor on other computers

After this lesson, you will be able to

- Monitor real-time performance data
- Record performance data into a log file
- Configure system and performance data alerts

Estimated lesson time: 20 minutes

Configuring System Monitor

With System Monitor, you can collect and view data by configuring counters that report hardware, application, and service activity for any computer on your network. Three configurations must be made for the data you wish to collect.

- **Type of data** You can specify one or more counter instances of performance monitor objects for which you want data to be reported.
- **Source of data** Either local or remote computer data can be collected by a counter. You must be a local administrator or a member of the Performance Log Users group on the computer from which you wish to collect data.
- **Sampling intervals** Data can be recorded manually in real time, or set to a periodic interval that you specify.

Viewing Data

When you first open System Monitor, three counters are loaded and begin to report real-time data:

- Memory: Pages/Second
- Physical Disk(_Total): Average Disk Queue Length
- Processor(_Total): % Processor Time

Figure 12-3 shows the System Monitor with the default counters loaded.

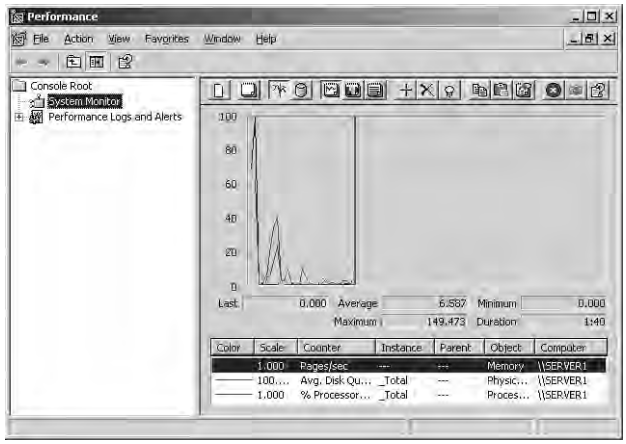

Figure 12-3 The System Monitor of the Performance console

Additional counters can be added or removed by choosing Add (Ctrl+I) on the toolbar, or right-clicking anywhere in the details pane and choosing Add Counters from the shortcut menu. In the Add Counters dialog box, you can select any of the available counters for either the local computer or any remote computer on your network. Counters are arranged and available for use based on the type of object, the counter in the object category, and the instance of the counter.

- **Object** A logical collection of resource, service, or application counters.
- **Counter** A data-reporting item. The data reported depends on the type of counter.
- **Instance** Refers to one or more occurrences of a counter, indexed by the number available on the computer. For example, on a computer with two processors, Instance "0" would refer to the first processor, Instance "1" to the second, and "_Total" the aggregate of both instances. In the case of a single instance of a counter, Instance "0" and " _Total" will be available.

Figure 12-4 shows the %Processor Time counter for Server01, which is a single-processor computer.

Figure 12-4 %Processor Time Counter–Single Processor

Exam Tip Remember that "_Total" represents the combined data from multiple instances of a counter when multiple instances are available.

Logging and Alerts

With Performance Logs And Alerts, you can collect performance data automatically from local or remote computers. You can view logged counter data by using System Monitor, or you can export the data to spreadsheet programs or databases for analysis and report generation. You can configure any counters available within System Monitor for use in Performance Logs And Alerts, with the following options:

- Collect data in a CSV or tab-separated format for exporting.
- View Counter Log data during logging and post-collection.
- Set Trace Logs (event-driven) based on available providers.
- Define parameters for the log file including start and stop times and maximum file size.
- Set an alert on a counter with options to send an administrative message, an application is executed, or a log is started when the configured threshold on the counter is breached.

Figure 12-5 shows the configuration dialog box for an alert on Server01 when Free Disk Space drops below 20 percent.

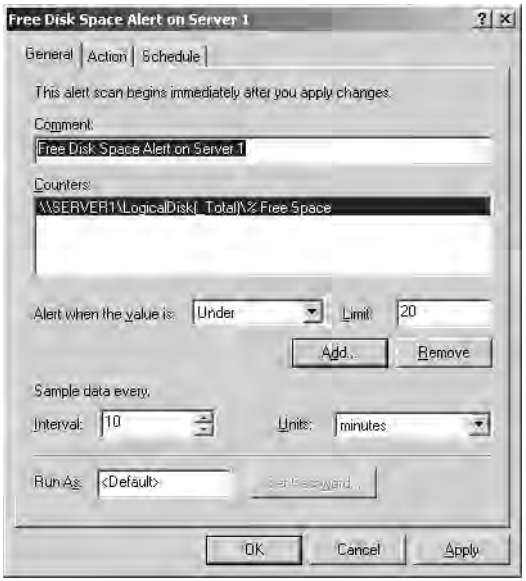

Figure 12-5 Alert Configuration for Disk Free Space

Real World **Monitoring Performance Data**

When monitoring performance data for your server or network, start from the top down; that is, start with the broadest monitoring configurations of % Processor Time, Disk and Processor Queue Length, Memory Use, and Network I/O to determine where the bottleneck occurs. Once you have determined the problem area, then look at the particular services and applications using the resource, and at protocol and thread levels, if needed. Usually, there is either one device or application causing the problem, or a global lack of resources on the system. Single devices can be reconfigured or replaced, and global resources can be added (more memory, faster processor, and so on) as appropriate.

The results of this monitoring can be ambiguous; however, if you do not have a baseline of system performance by which to judge your monitoring results. As soon as is practical after configuring a new computer, perform a set of monitoring activities for the key Processor, Memory, Network, and Process (Application and Services) objects to determine how your computer performs under normal conditions—commonly called a *baseline*—in normal, idle, and peak performance states. When problems or bottlenecks occur during later monitoring, measurement against the baseline will help to find a solution.

Decisions About Objects and Counters

The object counters that you choose in monitoring a server, either for a baseline or ongoing performance evaluation, can be considered in one of two ways. One method of server monitoring examines the role that the server performs in the environment and the corresponding demands placed on that server by the user population. Another view of server monitoring involves examining object categories of counters such as Processor, Memory, Network Interface, and PhysicalDisk, with less emphasis on the role that the server fulfills and more on a consistent monitoring standard.

Server Roles

Monitoring by server role is useful when servers perform within a single role in the network environment. These roles are defined by the services or resources that the server provides to the users. Examples of server roles include domain controllers, file servers, and Web servers. A server's demand for resources can be matched, in a performance monitoring situation, with the appropriate object counters that measure the resources most heavily used by a server in that role. Ongoing performance monitoring data can be compared to baseline data for optimization within that role. Table 12-1 outlines the objects that are commonly used when analyzing a server by its role.

Table 12-1 Server Roles and Objects To Be Monitored

Server role	Resources used	Objects and counters
Application servers	Memory, network, and processor cache	Memory, Processor, Network Interface, and System
Backup servers	Processor and network	System, Server, Processor, and Network Interface
Database servers	Disks, network, and processor	PhysicalDisk, LogicalDisk, Processor, Network Interface, and System
Domain controllers	Memory, processor, network, and disk	Memory, Processor, System, Network Interface, protocol objects (network-dependent, but can include TCPv4, UDPv4, ICMP, IPv4, NBT Connection, NWLink IPX, NWLink NetBIOS, and NWLink SPX), PhysicalDisk, and LogicalDisk
File and print servers	Memory, disk, and network components	Memory, Network Interface, PhysicalDisk, LogicalDisk, and Print Queue
Mail/messaging servers	Processor, disk, network, and memory	Memory, Cache, Processor, System, PhysicalDisk, Network Interface, and LogicalDisk
Web servers	Disk, cache, and network components	Cache, Network Interface, PhysicalDisk, and LogicalDisk

For each server role, create a baseline using the counters within each object appropriate for the role, and periodically examine each of the servers for significant changes.

Object Categories

In a network environment where servers perform within multiple roles, role-based monitoring can leave important gaps in monitored data. In such cases, more complete data should be collected from each of the primary object categories.

Memory Counters After you have established a baseline for memory use, periodic monitoring should be performed for deviations from that baseline. The following counters are useful in monitoring computer system memory:

- **Memory shortages: Memory\Available Bytes, Available Kbytes, or Available MBytes (to see the amount in megabytes); Process (All_processes)\Working Set; Memory\Pages/sec; Memory\Cache Bytes.** These counters show how much memory is taken up by all processes, and how much memory is available.
- **Frequent hard page faults: Memory\Pages/sec; Process (All_processes) \Working Set; Memory\Pages Input/sec; Memory\Pages Output /sec.** Hard page faults occur when a page of memory is needed but has been placed (swapped) into virtual memory. Excessive swapping degrades the performance of the computer, and can be addressed either by reducing the demands on the computer or increasing the amount of physical RAM.

Network Counters Network counters report data from the network interface cards (NICs) installed in the computer, and from the segment on which the NICs communicate. The following counters are useful in measuring the performance of a computer on the network:

- **Network Interface\Output Queue Length; Bytes Total\sec.** The Queue length should be low, and the total bytes high, which indicates a network card that is transferring packets quickly and without delay.
- **Network Interface: Bytes Sent/Sec; Current Bandwidth; Bytes Received/ Sec.** High values in these counters consistently and over time indicate that a network is being expected to carry more traffic than is optimal. Segmenting the network into smaller pieces or increasing the bandwidth of the network will decrease the chances of bottlenecks due to excessive traffic.

Note Different types of network configurations will allow for various levels of traffic efficiency and volume. When monitoring %Network Utilization, for example, 30 percent utilization is the maximum recommended for an unswitched Ethernet network. This means that a 10 megabyte (MB) Ethernet network becomes bottlenecked when its throughput exceeds 3 MB per second. If the value of the counter is above 40 percent, data collisions begin to hamper the performance of the network.

Process Counters For each demand on a system resource, there is often a process that is the instrument of that demand. Using process counters allows for viewing the individual processes (including system services) that are using system resources. The following are important counters to use when gathering process-based performance data:

- **Memory leaks; memory-intensive applications: Memory\Pool Nonpaged Allocs; Memory\Pool Nonpaged Bytes; Memory\Pool Paged Bytes; Process(process_name)\Pool Nonpaged Bytes; Process(process_name)\Handle Count; Process(process_name)\Pool Paged Bytes; Process(process_name)\Virtual Bytes; Process(process_name)\Private Bytes.** These counters show memory use by individual processes, allowing for redistribution of intensive applications (or isolation of applications with memory leaks) to other computers.

Note An application memory leak can be diagnosed by running that application on its own server, and monitoring for memory use that increases over time with no change in demand for services. This increase without a corresponding reason can indicate a memory leak.

Disk Counters The PhysicalDisk object counters provide data on activity for each of the hard disk storage devices, and the LogicalDisk object counters provide data on defined volumes (C:\, D:\, and so on) in your system. Monitoring LogicalDisk free space and PhysicalDisk performance counters will provide useful data. The following are important counters for Physical and Logical Disk monitoring:

- **LogicalDisk\% Free Space.** This counter reports the percentage of unallocated disk space to the total usable space on the logical volume. This counter is not available for a physical disk.

Note When calculating the _Total instance, the %Free Space counters recalculate the sum as a percentage for each disk.

- **PhysicalDisk\Avg. Disk Bytes/Transfer; \Avg. Disk sec/Transfer; \Avg. Disk Queue Length; \% Disk Time.** These counters measure the size of input/output (I/O) operations over time, and how busy the drive is, performing the requested disk activity. The disk is efficient if it transfers large amounts of data relatively quickly, and has a queue length <2 over time for each disk spindle.

Practice: Using the Performance Console

In this practice, you will record Performance data, analyze the data in System Monitor, and export the data for import into an Excel spreadsheet.

Exercise 1: Recording Performance Data

In this exercise, you will create a log file with LogicalDisk, PhysicalDisk, and Server Work Queue data.

1. Log on to Server01 as an administrator, and start the Performance console.
2. Expand Performance Logs And Alerts in the folder pane, and then select Counter Logs.
3. In the detail pane, right-click and select New Log Settings from the shortcut menu.
4. Create a log file called Test, and add the LogicalDisk, PhysicalDisk, and Server Work Queues objects to the log, and set the data sampling interval to 8 seconds. Take note of the file name and location for the log, and then click OK to start the log.
5. As the log is recording, perform some activities with other applications on your computer. After approximately 30 seconds, return to Performance Logs And Alerts and stop the log recording.
6. In System Monitor, click View Log Data (Ctrl+L, or fourth button from the left), and load the log file from your test.

The graph in System Monitor now shows the recorded data from your logging session. You can now change the views between graph, histogram, and report to see the data in different ways. This log file that you have created is in the default format (Binary File), strictly for use in the Performance Console.

Exercise 2: Importing Logged Data

In this exercise, you will save the logged data from Exercise 1 for import into Microsoft Excel.

1. If needed, reopen the Performance console.
2. Right-click the Test log file setting, and then choose Properties.
3. In the Test Properties dialog box, click the Log Files tab, and then change the Log File Type from Binary File to Text File (Comma Delimited).
4. Click OK, and then start the log file recording. Perform some disk-related tasks on your computer for approximately 30 seconds, and then stop the log recording.

The log file you have created is in CSV format, and can be opened, viewed and analyzed in Excel.

Note If you intend to load the CSV file into Excel, Performance Logs And Alerts cannot have the file open because Excel requires exclusive access to the file to open it.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. Your goal is to monitor all your Windows Server 2003 servers so that they can be defragmented on a regular schedule, and as efficiently as possible. The disk defragmentation program that you use requires at least 20% free disk space on each volume in order to defragment properly. What should you do?

2. You have been monitoring one of your Windows Server 2003 servers due to poor performance on the network. The following data is representative of your findings:

   - ❑ Processor: % Processor Time: High
   - ❑ Physical Disk: % Disk Time: Low
   - ❑ Memory: Pages/sec: Low
   - ❑ Processor: Interrupts/sec: High
   - ❑ Process: % Processor Time (for non-service processes): Low
   - ❑ Process: % Processor Time (for system services): Low

   What is the most likely explanation for the problem?

3. The server that you are using to monitor the other servers on your network is overburdened with the task, so you must lighten its load of monitoring. To make the greatest impact for the monitoring computer's performance while maintaining as much monitored data as possible, what should you do?

Lesson Summary

The Performance console has two snap-ins configured: System Monitor and Performance Logs And Alerts. The System Monitor is designed for real-time reporting of data to a console interface, and can be reported in graph, histogram, or numeric form. The Performance Logs And Alerts snap-in is designed to write data to a file (log) and report counter values that breach a threshold (alert). Logs written by Performance Logs And Alerts can be loaded into System Monitor for analysis, and exported to various file types (such as CSV and HTML) for reporting purposes.

Lesson 3: Using Task Manager

Task Manager provides information about programs and processes running on your computer. It also displays several common process performance counters.

After this lesson, you will be able to

- Configure Task Manager to display performance data
- Use Task Manager to start and end applications and processes

Estimated lesson time: 15 minutes

Task Manager Overview

Task Manager can be opened by right-clicking an open area of the taskbar and then choosing Task Manager, or by pressing CTRL+ALT+DEL and then choosing Task Manager. The Windows Server 2003 Task Manager interface, by default, presents five tabs on which its data is categorized: Applications, Processes, Performance, Networking, and Users.

Applications Tab

The Applications tab shows the status of the user-level programs running on the computer. Services and system applications running in a context different from the logged on user are not displayed. On the Applications tab you can also start a new program with New Task, end a program with End Task, or switch to another program using Switch To. By right-clicking on an application, you can also select Go To Process from the shortcut menu, which will take you to the corresponding process on the Process tab. Figure 12-6 shows the Applications tab of Task Manager.

Figure 12-6 The Applications tab of Task Manager

Processes Tab

The Processes tab shows information about all processes running on your computer, including user-level applications, services, and other system processes. By choosing Select Columns from the View menu, you can add or remove columns of data including memory usage changes (deltas), process IDs, and processor use. You can sort by any column by clicking on the column header.

By right-clicking on any process, you can change the priority of processor time that the process receives, set the processor affinity on multiple processor computers, and end a process. For processes that have child or related processes, you can end all related processes by choosing End Process Tree. If you needed to end a mail application, for example, you might also need to end the MAPI spooler; appropriately, you would right-click on the mail application and choose to End Process Tree. The Processes tab is shown in Figure 12-7.

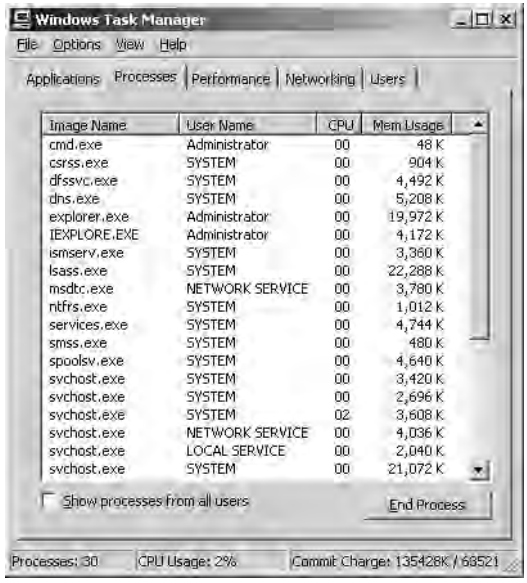

Figure 12-7 The Processes tab of Task Manager

Caution Changing settings of a process such as priority or processor affinity can have an adverse effect on the performance of other applications running on your computer. Ending a process, especially a process tree, should be done only after normal termination procedures have failed. Windows Server 2003, thankfully, safeguards its processes from termination through Task Manager, but they are still susceptible to resource starvation through inappropriate priority adjustment of other processes.

Performance Tab

The Performance tab displays a real-time view of key elements of your computer's performance. Graphs are presented for each processor on the system and memory usage. Text displays show physical, kernel, and commit memory; also, the number of handles and threads in use by active processes are displayed. The Performance tab is shown in Figure 12-8.

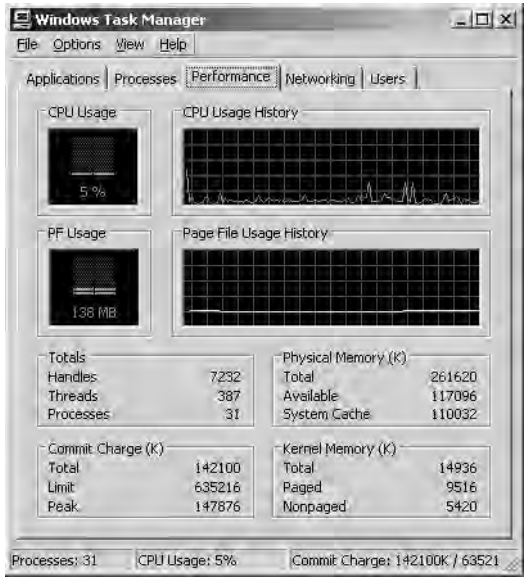

Figure 12-8 The Performance tab of Task Manager

Networking Tab

The Networking tab shows all active network connections by name, their connection speed, bandwidth usage, and status. The Networking tab is shown in Figure 12-9.

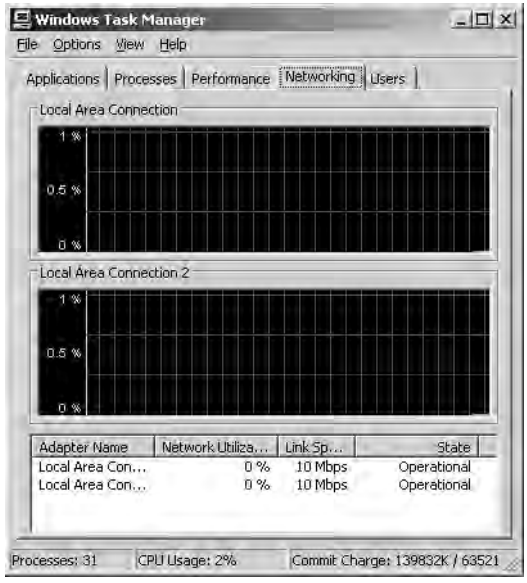

Figure 12-9 The Networking tab of Task Manager

Users Tab

The Users tab shows all users who are logged on, and allows for the logoff or forced disconnection of the user from the computer. Logged-on users may be local at the console, or remotely attached from the network. Network messages can be sent to remote users (it certainly is polite to tell them before you disconnect them) by selecting the users' session and then clicking Send Message. The Users tab is shown in Figure 12-10.

Figure 12-10 The Users tab of Task Manager

Practice: Task Manager

In this practice, you will use Task Manager to start an application and identify its process.

1. Right-click an open section of the taskbar, and choose Task Manager from the shortcut menu.
2. On the Applications tab, click New Task. Type **explorer**, and then click OK.

   Windows Explorer will open focused on its default window (typically My Documents), and the My Documents (or other focus window) application name will appear in the Application tab of Task Manager.
3. Right-click the newly-opened application name, and choose Go To Process.

The focus in Task Manager changes to the Processes tab, and the Explorer.exe process is highlighted. From this point, as a situation involving applications and Task Manager warrants, you can adjust the priority of the process or end it.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. What information can Task Manager provide about the performance of applications?

   _____

   _____

2. Your computer crashes with almost clocklike predictability approximately one hour after each system startup. You suspect an application with a memory leak that is causing the system to run out of memory. How can you use Task Manager to determine which application is causing the problem?

   _____

   _____

3. You are running a database application on your computer. Your computer has two processors. You want the database application to run on the second processor. How can you use Task Manager to do this?

   _____

   _____

Lesson Summary

Task Manager provides dynamic views into current performance of your computer as it relates to running processes and applications. With configurable refresh intervals and selectable columns of data, the Task Manager shows the processor, memory, and I/O usage by processes. Applications can be started or ended from the Applications tab, and processes can be elevated in priority or terminated, including child processes, from the Processes tab. The Performance tab gives an aggregate view of processor and memory use on the computer. The Networking tab does the same for network utilization and basic configuration data. The Users tab, if available, will allow for logoff of a local session or disconnection of a remote session. Remote sessions can also have messages sent to the connected user.

Lesson 4: Using the WMI Event Logging Provider

Windows Management Instrumentation (WMI) is the Microsoft implementation of Web-Based Enterprise Management (WBEM), an initiative to establish standards for creating, reading, and modifying management information. WMI is WBEM-compliant and provides integrated support for the Common Information Model (CIM), the data model that describes the objects that exist in a management environment. The WMI repository is the database of object definitions, and the WMI Object Manager handles the objects as input from WMI providers. The WMI providers can receive a wide variety of input from services, applications, and system components.

After this lesson, you will be able to

- Use WMI and WMI command-line (WMIC) to monitor running services
- Use WMI and WMIC to identify installed programs
- Use WMI and WMIC to report Event data

Estimated lesson time: 30 minutes

How WMI Works

Described briefly, WMI sources of information ("providers") output information about their components (devices, services, applications, and so on) to the WMI Object Manager, which enters the information into the WMI database ("repository"). Depending on what is accepted as input and returned as output by each provider, administrators will be able to use methods to manipulate the components, set properties, and configure events that can alert administrators to changes in the components. The WMI Repository can be accessed by management tools supplied by a system, application, or device vendor; through Application Programming Interfaces (APIs) or scripting tools (Windows Scripting Host, for example), or from the command line using Windows Management Interface Command-line (WMIC).

Off the Record You do not have to become a scripting expert to leverage WMI, and you have been using WMI even if you did not realize it. Many tools already leverage WMI to report and configure the object of the WMI provider. For example, several Windows Server 2003 tools that use WMI are System Information, System Properties, and Services.

Windows Management Interface Command-line (WMIC)

The WMIC provides a command-line interface to WMI, and can be used to manage, locally or remotely, any computer with WMI that can authenticate the user running WMIC. For WMIC to manage a remote computer, only WMI needs to be on the local

computer from which the monitoring activity will be accomplished; WMIC does not have to be available on the remotely managed computer. You can use WMIC to accomplish various types of tasks:

- **Local management of a computer** You are at the computer and use the WMIC command to manage it.
- **Remote management of a computer** You are at the console of one computer and use WMIC to manage another computer.
- **Remote management of multiple computers** You are at the console of one computer and use WMIC to manage multiple computers with a single command.
- **Administrative scripting** You use WMIC to write a management script (batch file) to automate the management of a computer (local, remote, or multiple computers).

Administration with WMIC

Although a complete discussion of WMIC use with WMI is beyond the scope of this chapter (there are several good books devoted to the subject), a few points of reference are needed so that you can answer questions about monitoring with WMI and WMIC effectively. Unless otherwise noted, the remainder of discussion points in this chapter assumes that you are using WMIC in interactive mode, giving you the ability to issue single commands and view the results from within the WMIC environment.

Exam Tip Any references to WMIC in interactive mode versus non-interactive mode have no effect on the way commands are structured or used. The difference between interactive and non-interactive modes has to do with how many commands you intend to execute, and whether these commands are being entered manually or in a batch file. Enter the interactive mode of WMIC by typing **wmic** at a command line, pressing Enter, and then typing **exit** or **quit** to leave. Non-interactive mode consists of a single-line command beginning with WMIC either at a command line or in a batch file.

WMI works within the context of a namespace, the default being root\cli (MSFT_cli in the XSL stylesheet) that controls what properties, methods (verbs), and aliases are available in WMI. You can add aliases, methods, and properties if necessary (put on your programming hat), but the list is robust enough for most monitoring tasks.

Security for WMI is configured through the WMI Control snap-in (Wmimgmt.msc), in the WMI MMC. By default, users have permissions to read WMI provider information through WMIC on a local computer, but do not have permission to connect remotely or write information outside of the provider context. Administrators who want to grant additional permissions to a user or group must do so through the WMI Control snap-in.

WMIC Aliases The first parameter of a WMIC command line is the alias. The alias name must be unique in the WMI namespace schema, and provides access to WMI information without needing to remember more complex schema objects and properties. Table 12-2 lists the properties associated with each alias instance. The complete alias and namespace lists, and other detailed information about WMIC aliases, can be found in Windows Server 2003 Help and Support: Alias Namespaces and Classes.

Table 12-2 WMIC Aliases

Property	Description
FriendlyName	The name of the alias; it must be unique.
Description	A description of the alias. This is the descriptive text when /? is entered at the WMIC command line.
Formats	A list, each of which has a name and a list of properties (objects of the class MSFT_CliProperty) to be displayed for that format. All formats are objects of the class MSFT_CliFormat.
Verbs	A list, each of which are the various behaviors available through this alias. The behaviors come in two forms: Standard verbs, which are directly supported by the utility. User-defined verbs, which must map to some method defined for the target of the alias. All verbs are objects of the class MSFT_CliVerb.
Qualifiers	A list, similar to WMI qualifiers. All qualifiers are objects of the class MSFT_Qualifier.
Target	A list, similar to WMI qualifiers. All qualifiers are objects of the class MSFT_Qualifier.
PWhere clause	Optional WHERE clause that limits the Target. It has substitution values which are the parameters of the alias. The substitution values are marked with #. If multiple parameters are needed, they are matched with the # markers in sequence.
Connection	Details on which computers to connect to, the security details to be used, and so on. If a connection is not specified, the computers to be accessed are the value of /NODE, and the namespace is the value of /NAMESPACE. If a user name and password are not provided, then the value of /USER and /PASSWORD, if available, is used (otherwise the current account is used).
View an alias schema	Use a metaalias Alias to view an alias schema; example: ALIAS OS.

WMIC Verbs Most aliases have actions that they perform: these actions are initiated by issuing a verb along with the alias. In combination with parameters and switches, these alias-verb combinations control what configuration is set within the application or system, or what information is read from the WMI Repository. Table 12-3 lists the key verbs used in monitoring and their descriptions. The complete verbs list, and other

detailed information about WMIC verbs, can be found in Windows Server 2003 Help and Support: WMIC Verbs.

Table 12-3 WMIC Verbs

Verb	Action	Parameters	Example
CALL	Executes methods	Method and parameter list if appropriate. Parameter lists are comma delimited. Use SERVICE CALL /? to get a list of available methods and their parameters for the current alias.	SERVICE WHERE CAPTION='TELNET' CALL STARTSERVICE
GET	Get specific properties	Property name or switch	PROCESS GET NAME
LIST	Show data	LIST is the default verb. There are many switches and adverbs that can be used with the LIST verb (example: BRIEF)	PROCESS LIST BRIEF

Using WMIC in Monitoring

With WMI running on a computer, and sufficient administrative credentials owned by the user running WMIC, local or remote monitoring of a computer is available at the command line. In non-interactive mode, multiple commands can be contained in a batch file that is run either manually or on an automated schedule. These WMIC commands can be output to a CSV file, text file, or HTML page to be viewed and analyzed. Following are examples of common monitoring scenarios and output that illustrate the use of WMIC for monitoring.

- PRODUCT

  This command will output to the console the results of a query for all installed software on the local computer.

- /OUTPUT:c:\applog.htm NTEVENT WHERE "eventtype<3 AND logfile='Application'" GET Logfile, SourceName, Eventtype, Message, TimeGenerated /FORMAT:htable:"sortby=EventType"

  This command will output to an HTML file (C:\applog.htm) any events with types 0, 1, or 2 from the Application Log of the local computer. The list will be formatted in an HTML table (using the XML stylesheet htable.xsl) and sorted by Event Type.

- /OUTPUT:c:\applog.csv /NODE:@"c:\serverlist.txt" NTEVENT WHERE "eventtype<3 AND logfile='Application'" GET Logfile, SourceName, Eventtype, Message, TimeGenerated /FORMAT:csv:"sortby=EventType"

  This command will output to a CSV file (c:\applog.csv) any events with types 0, 1, or 2 from the Application Logs of the computers in the file serverlist.txt. The list will be formatted in a comma-separated list (using the XML stylesheet csv.xsl) and sorted by Event Type.

- OS ASSOC

  This command displays information related to the operating system hotifixes and patches that have been installed.

Practice: WMI Data from Event Viewer

In this practice, you will extract data from the Event Viewer and publish it to a Web page.

1. Logged on as Administrator on Server01, open a command window, type **wmic** and press Enter. This enters WMIC in interactive mode.
2. At the WMIC prompt, type the following command to access the Security Log data from Lesson 1, Exercise 3:

   NTEVENT WHERE "EVENTTYPE=5 AND LOGFILE='SECURITY'" GET LOGFILE, SOURCENAME, EVENTTYPE, MESSAGE, TIMEGENERATED

   This outputs the Failure Audit entries to the console.
3. At the WMIC prompt, type the following command to output the same information to a Web page called C:\seclog.htm.

   /OUTPUT:C:\seclog.htm NTEVENT WHERE "EVENTTYPE=5 AND LOGFILE='SECURITY'" GET LOGFILE, SOURCENAME, EVENTTYPE, MESSAGE, TIMEGENERATED /FORMAT:htable
4. Double-click the file C:\Seclog.htm to open the file in Internet Explorer.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You need to get patch and hotfix information from a number of servers on your network. You would like to do this remotely. How can you use WMI to accomplish the task?

   _____

   _____

2. You want to get a list of all installed applications on 17 computers in the development department. You would like to do this remotely. How can you use WMI to accomplish this?

   _____

   _____

3. You want to give a small group of engineers the ability to use WMI to get information from some of the development servers, but you do not want to give them administrator privileges on the servers. What can you do to give the engineers access?

Lesson Summary

WMI is a WBEM-compliant utility that uses a CIM-compliant database of management information collected by running on each Windows Server 2003 computer. The command line interface for WMI is WMIC, which uses a series of aliases, verbs, switches and parameters to change configuration on or get information from a computer system. WMIC can connect to any computer remotely, so long as the user initiating the connection has sufficient privileges on the remote computer. The local administrator on a computer has permission to connect remotely, so Domain Administrators each have the ability to perform remote administration with WMI and WMIC. For archiving and reporting purposes, WMI data can be output through WMIC to CSV or HTML pages. Multiple computers can have commands issued to them either from the command line or from a text file. With the exception of needing to include the WMIC command at the beginning of each line, issuing commands from a batch file in non-interactive mode is no different from using WMIC in interactive mode.

Case Scenario Exercise

You have been placed in charge of the Information Technology department at your new company, and are trying to put better practices in place than your predecessor did. Server hardware failures are quite commonplace, and user satisfaction with network performance is very low. From your initial interviews with your administrative staff, there has been very little planning and very much fire-fighting over the past several years.

The mandate from the company's vice president is that you get things cleaned up quickly, and relieve his anxiety that important company data is walking out the door due to lax security practices. User productivity needs to improve, and that includes all the applications that the users have installed on their computer systems—against company policy.

Here, then, is the checklist you create for improving the technology environment:

- Use WMI and other file utilities to get a complete list of all file-based resources that are located in servers on the network. Completely document the permission structure on those resources and cross-reference it with the department heads to be certain that you understand the file resource access needs clearly.

- Use Event Viewer and the Performance console to get an accurate picture of any immediate bottleneck problems due to device failure, service misconfiguration, or application incompatibilities. Replace hardware, properly configure services, and upgrade applications where necessary to improve the component parts of the running environment.
- Once the permissions are defined, put Failure Access Auditing in place to find anyone who is attempting to gain unauthorized resource access, and through what means.
- Use Performance Logs And Alerts to baseline the servers once clearly defined bottlenecks have been removed. Continue to monitor for changes in server performance against the baseline.

Troubleshooting Lab

Users in the Help Desk group have been creating their own Web pages to publish technical data for the rest of the group, and have many utilities that they use periodically in testing applications for functionality and stability. Recently, these users have been asking for some help in determining why their computers' performance has recently declined significantly.

Using the Performance console, take a baseline of the following counters:

- Cache\Data Map Hits %
- Cache\Fast Reads/sec
- Cache\Lazy Write Pages/sec
- Logical Disk\% Free Space
- Memory\Available Bytes
- Memory\ Pool Nonpaged Allocs
- Memory\ Pool Nonpaged Bytes
- Memory\ Pool Paged Allocs
- Memory\ Pool Paged Bytes
- Processor(_Total)\% Processor Time
- System\Context Switches/sec
- System\Processor Queue Length
- Processor(_Total)\Interrupts/sec

Monitor each of the suspect computers for one week of normal activity, recording the resulting output in a log file unique to each computer. Use a remote computer to collect the monitoring data so as not to skew the results of your baseline.

Analyze the data to determine if there are any obvious bottlenecks. This list of counters is particularly baselining memory, disk I/O, and processor performance on each of the computers. Once the bottleneck has been defined, the applications (processes) should be examined to determine which of them are the heaviest contributors to the problem. The applications can then be upgraded, if that helps; removed, or resources can be added to the computers sufficient to perform the required tasks.

Chapter Summary

- Event Viewer presents data in the form of logs. The Application, System, and Security logs are on every Windows Server 2003 server. Domain controllers have two additional logs relating to Active Directory, and other application servers (such as DNS) have their own set of log files.
- The Performance console (perfmon.msc) consists of two snap-ins: System Monitor and Performance Logs And Alerts. System Monitor shows real-time performance data based on Object counters, and can display the log data recorded by Performance Logs And Alerts either in the form of Counter (interval polling) logs, or Trace (event-driven) logs.
- Task Manager is used to view real-time performance data surrounding processes and applications. Processes can be initiated and ended using Task Manager. Processes can also be adjusted up or down in CPU priority, and can be assigned affinity to a particular processor on a multiprocessor computer.
- WMI is a management system that collects data from computer systems. The control interface of WMI Control snap-in allows for adjustment of permissions beyond the default of the local administrator to manage computers across the network. While WMI is capable of configuring many different types of system behavior including users, groups, and services, the focus of this chapter is on the ability to extract data from the WMI Repository using the command line interface to WMI, WMIC. WMIC is capable of reporting running services, installed applications, and publishing Event Viewer data to CSV or HTML files for ease of distribution and analysis.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- Event Viewer does not perform configuration, but collects data from different reporting providers. Data reported is organized into the appropriate log, and can be filtered, sorted, and exported for ease of analysis.
- Task Manager is a tool used only on the local computer, and does not allow configuration of memory, processor, or other settings. Task Manager is exclusively used to start, stop, prioritize, and set processor affinity for applications.
- The Performance Logs And Alerts snap-in can do no configuration, only reporting data through Counter Logs as reported by providers (object counters) on a configured interval, or through Trace Logs as reported by event-driven providers.
- WMI requires administrative credentials for access to the remote computer for configuration of settings.
- WMIC is not an Active Directory Schema Management Tool. WMI maintains its own schema.

Key Terms

Windows Management Instrumentation (WMI) The Microsoft implementation of Web-Based Enterprise Management Initiative to establish standards of data in Enterprise Management

Windows Management Instrumentation Control (WMIC) A command line utility that interfaces with the WMI Repository (database) for configuration and monitoring management

Task Manager An interface tool for the manipulation of processes

System Monitor A component of the Performance console, as is the Performance Logs And Alerts snap-in, and should not be confused with System Properties

13 Recovering from System Failure

Exam Objectives in this Chapter:

- Perform Automated System Recovery (ASR)
- Perform server system recovery

Why This Chapter Matters

Although Microsoft Windows Server 2003 offers superior levels of stability and reliability, power supplies, cooling fans, chip sets and yes, even code, can cause a computer to fail. And when a server fails in the forest, everyone hears it fall. Throughout this training kit, you have learned how to implement and support best practices that will minimize the risk of failure. You have also learned how to recover from the failure of specific services, drivers, and hardware configurations. In this chapter, you will learn the remaining skills that are required to recover a server when the operating system itself is corrupted or inaccessible due to catastrophic failure.

Lessons in this Chapter:

Before You Begin

This chapter covers the concepts and skills related to recovering a failed server. To complete the exercises in this chapter, prepare the following:

- A computer running Windows Server 2003. The examples use the computer name Server01. It can be a member server or a domain controller. Backups that are created during the exercises will complete more quickly if the computer is a member server.
- A second physical disk is required to perform the exercise that demonstrates Automated System Recovery.
- If you complete the Automated System Recovery exercise, all data on the disk containing the system volume will be erased. Do not perform the Automated System Recovery if you want to maintain any data on that disk.

Lesson 1: Recovering from System Failure

In a worst-case scenario, server hardware fails and cannot be recovered. To return to operations, you must have a complete backup of the server that you can restore to a new piece of hardware. This complete backup will include data stored on the server, applications, and the operating system itself. In Chapter 7, you learned how to use the Backup Utility and the Ntbackup command-line tool to back up data. In this lesson, you will learn how to use the same utilities to back up the system so that you can return to operational status quickly in the event of such a worst-case scenario. You will also learn how to use the Recovery Console to perform surgical repairs of specific problems including service or driver failures.

After this lesson, you will be able to

- Back up the System State
- Prepare an ASR backup set and repair a computer using Automated System Recovery
- Install and use the Windows Server 2003 Recovery Console

Estimated lesson time: 60 minutes

A Review of Recovery Options

Throughout this book, we have addressed methods used to repair and recover from specific types of failures:

- Data loss or corruption: Chapter 7 discussed the backup and restore of data as well as the Volume Shadow Copy Service, the new feature in Windows Server 2003 that allows users to access or restore previous versions of files in shared folders on servers.
- Driver updates resulting in system instability: Chapter 10 introduced the new driver rollback capability of Windows Server 2003. If a driver has been updated and the system becomes unstable, that driver and any new settings that were configured can be rolled back to a previously installed version and state. Printer drivers cannot be rolled back. You also learned that it is easy, using Device Manager, to disable a device that causes instability. If an application or supporting software contributes to the instability, use Add Or Remove Programs to remove the offending component.
- Driver or service installation or update results in the inability to start the system: Chapter 10 covered the use of the Last Known Good Configuration, which rolls back the active ControlSet of the system's registry to the ControlSet that was used

the last time a user successfully logged on to the system. If you install or update a service or driver and the system crashes or cannot reboot to the logon screen, the Last Known Good Configuration effectively takes you back to the version of the registry that was active before the driver or service was installed. You also learned about the variety of Safe mode options, which enable the system to start with specific drivers or services disabled. Safe mode can often allow you to start an otherwise unbootable computer and, using Device Manager, disable, uninstall, or roll back a troublesome driver or service.

- Failure of the disk subsystem: Chapter 11 discussed the steps required to configure disk redundancy through mirrored (RAID-1) or RAID-5 volumes, and how to recover from the failure of a single disk within a fault-tolerant volume.

Each of these recovery and repair processes makes the assumption that a system can be restarted to some extent. When a system cannot be restarted, the System State, Automated System Recovery, and the Recovery Console can return the system to operational status.

System State

Windows 2000 and Windows Server 2003 introduced the concept of *System State* to the backup process. System State data contains critical elements of a system's configuration including:

- The system's registry
- The COM+ Class Registration Database
- The boot files, which include boot.ini, ntdetect.com, ntldr, bootsect.dos, and ntbootdd.sys
- System files that are protected by the Windows File Protection service

In addition, the following are included in the System State when the corresponding services have been installed on the system:

- Certificate Services database on a certificate server
- Active Directory and the Sysvol folder on a domain controller
- Cluster service information on a cluster server
- Internet Information Services (IIS) metabase on a server with IIS installed

To back up the System State in the Backup Utility, include the System State node as part of the backup selection. The System State and its components are shown in Figure 13-1.

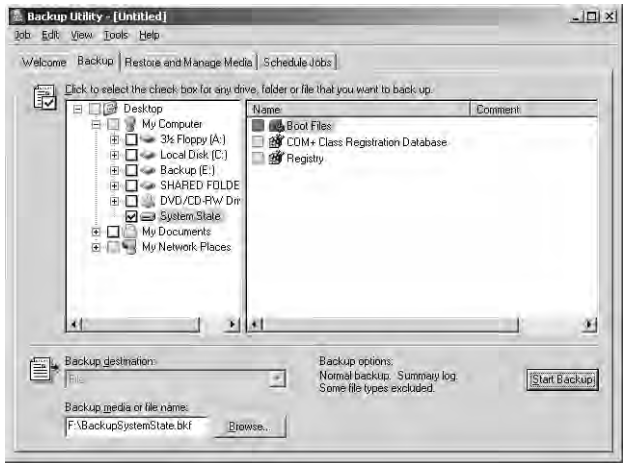

Figure 13-1 The System State

If you prefer to use the command line, use Ntbackup with the following syntax:

```
Ntbackup backup systemstate /J "backup job name" ...
```

Followed by the /F switch to indicate backing up to a file, or appropriate /T, /G, /N, /P switches to back up to a tape. The switches for the Ntbackup command are described fully in Chapter 7.

There are several important notes and considerations related to backing up the System State:

- You cannot back up individual components of the System State. For example, you cannot back up the COM+ Class Registration Database alone. Because of interdependencies among System State components, you can back up only the collection of System State components as a whole.
- You cannot use Ntbackup or the Backup Utility to back up the System State from a remote machine. You must run Ntbackup or the Backup Utility on the system that is being backed up. You can, however, direct the backup to a file on a remote server, which can then transfer the file onto another backup media. Or you can purchase a third-party backup utility that can remotely back up the System State.
- The System State contains most elements of a system's configuration, but may not include every element required to return the system to full operational capacity. It is therefore recommended to back up all boot, system, data, and application volumes when you back up the system state. The System State is a critical piece of a complete backup, but is only one piece.

- Performing a system state backup automatically forces the backup type to Copy, although the interface may not indicate that fact. Take that fact into consideration when planning whether to include other items in your backup selection.

To restore the System State on a computer that is operational, use the Backup Utility and, on the Restore And Manage Media tab, click the System State check box. If the computer is not operational, you will most likely turn to Automated System Recovery to regain operational status.

System State on a Domain Controller

The System State on a domain controller includes the Microsoft Active Directory directory service and the Sysvol folder. You can back up the System State on a domain controller just as on any other system, using the Backup Utility or Ntbackup command. As with all backup media, it is paramount to maintain physical security of the media to which the Active Directory is backed up.

To restore the System State on a domain controller, you must restart the computer, press F8 to select startup options, and select Directory Services Restore Mode. This mode is a variation of the Safe modes described in Chapter 10. In Directory Services Restore Mode, the domain controller boots but does not start Active Directory services. You can log on to the computer only as the local Administrator, using the Directory Services Restore Mode password that was specified when Dcpromo was used to promote the server to a domain controller.

When in Directory Services Restore Mode, the domain controller does not perform authentication or Active Directory replication, and the Active Directory database and supporting files are not subject to file locks. You can therefore restore the System State using the Backup Utility.

When restoring the System State on a domain controller, you must choose whether to perform a non-authoritative (normal) or authoritative restore of the Active Directory and Sysvol folder. After restoring the System State using the Backup Utility, you complete a non-authoritative restore by restarting the domain controller into normal operational status. Because older data was restored, the domain controller must update its replica of the Active Directory and Sysvol, which it does automatically through standard replication mechanisms from its replication partners.

There may be occasions, however, when you do not want the restored domain controller to become consistent with other functioning domain controllers and instead want all domain controllers to have the same state as the restored replica. If, for example, objects have been deleted from Active Directory, you can restore one domain controller

with a backup set that was created prior to the deletion of the objects. You must then perform an authoritative restore, which marks selected objects as authoritative and causes those objects to be replicated *from* the restored domain controllers *to* its replication partners.

To perform an authoritative restore, you must first perform a non-authoritative restore by using the Backup Utility to restore the System State onto the domain controller. When the restore is completed and you click Close in the Backup Utility, you are prompted to restart the computer. When that occurs, you must select No. Do not allow the domain controller to restart. Then, open a command prompt and use Ntdsutil to mark the entire restored database or selected objects as authoritative. You can get more information about Ntdsutil and authoritative restore by typing **ntdsutil /?** at the command prompt or by using the online references in the Help And Support Center. The *MCSE Training Kit (Exam 70-294): Planning, Implementing, and Maintaining a Microsoft Windows Server 2003 Active Directory Infrastructure* (Microsoft Press, 2003) addresses domain controller recovery in detail.

Exam Tip What is most important to remember for the 70-290 exam is that the System State can only be restored on a domain controller by restarting the domain controller in Directory Services Restore Mode, and that Ntdsutil is used to recover deleted objects in Active Directory by marking those objects as authoritative, following a normal, or non-authoritative, restore of the System State with the Backup Utility.

Automated System Recovery

Recovering a failed server has traditionally been a tedious task, involving reinstallation of the operating system, mounting and cataloging the backup tape, then performing a full restore. Automated System Recovery makes that process significantly easier. Automated System Recovery requires you to create an ASR set, consisting of a backup of critical system files, including the registry, and a floppy disk listing the Windows system files that are installed on the computer. If the server ever fails, you simply restart with the Windows Server 2003 CD-ROM and select the option to perform an Automated System Recovery. The process uses the list of files on the ASR disk to restore standard drivers and files from the original Widows Server 2003 CD-ROM, and will restore remaining files from the ASR backup set.

To create an ASR set, open the Backup Utility from the Accessories program group, or by clicking Start, then Run, and typing **Ntbackup.exe**. If the Backup And Restore Wizard appears, click Advanced Mode. Then, from the Backup Utility's Welcome tab, or from the Tools menu, select ASR Wizard. Follow the instructions of the Automated

System Recovery Preparation Wizard. It will request a 1.44 megabyte (MB) floppy disk to create the ASR floppy. The ASR Wizard is shown in Figure 13-2.

Figure 13-2 The Backup Destination page of the ASR Wizard

The backup created by the ASR Wizard includes disk configuration information for each disk in the computer, a System State backup, and a backup of files including the driver cache. The backup set is sizable. On a standard installation of Windows Server 2003, the ASR backup size will be more than 1 gigabyte (GB).

The ASR floppy disk is created by the Automated System Recovery Preparation Wizard, and is specific to the system and the time at which the ASR set was created. You should label the ASR backup set and floppy disk carefully and keep them together.

The ASR floppy disk contains two catalogs of files on the system: Asr.sif and Asrpnp.sif. If the system does not have a floppy drive when you create the ASR set, you can create the floppy disk after running the wizard by copying these two files from the %*Systemroot*%\repair folder on the system to another computer that does have a floppy drive, and copying the files to the floppy disk on that second system. If you lose the floppy disk, you can restore the two files from the %*Systemroot*%\repair folder in the ASR backup set. You *must* have the ASR floppy disk to perform an Automated System Recovery. If the system does not have a floppy drive you will need to connect one before performing the restore.

Tip The ASR set contains the files required to start the system. It is not a comprehensive backup of the entire system. Therefore it is highly recommended to create a complete backup, including the System State, system volume, applications and, perhaps, user data when you create your ASR set.

When you perform an Automated System Recovery, you will need

- The Windows Server 2003 setup CD-ROM
- The ASR backup set
- The ASR floppy disk created at the same time as the ASR backup set

Tip You will also need any mass storage device drivers that are not part of the standard Windows Server 2003 driver set. To facilitate recovery, you should consider copying those drivers to the ASR floppy disk.

To restore a system using Automated System Recovery, restart using the Windows Server 2003 CD-ROM, just as if you were installing the operating system on the computer. If the computer requires a mass storage device driver that is not included with Windows Server 2003, press F6 when prompted and provide the driver on a floppy disk. After loading initial drivers, the system will prompt you to press F2 to perform an Automated System Recovery. Press F2 and follow the instructions on your screen. Automated System Recover will prompt you for the system's ASR floppy, which contains two catalogs, or lists, of files required to start the system. Those files will be loaded from the CD-ROM. Automated System Recovery will restore remaining critical files, including the system's registry, from the system's ASR backup set. There is a restart during the process, and if the computer requires a vendor-specific mass storage device driver, you will need to press F6 during this second restart as well. Because there is a restart, you should either remove the floppy after the initial text-based portion of the restore, or set the restart order so that the system does not attempt to restart from the floppy drive.

Recovery Console

The Recovery Console is a text-mode command interpreter that allows you to access to the hard disk of a computer running Windows Server 2003 for basic troubleshooting and system maintenance. It is particularly useful when the operating system cannot be started, as the Recovery Console can be used to run diagnostics, disable drivers and services, replace files, and perform other targeted recovery procedures.

Installing the Recovery Console

You can start the Recovery Console by booting with the Windows Server 2003 CD-ROM and, when prompted, pressing R to choose the repair and recover option. However, when a system is down you will typically want to recover the system as quickly as possible, and you may not want to waste time hunting down a copy of the CD-ROM or waiting for the laboriously long restart process. Therefore, it is recommended to proactively install the Recovery Console.

To install the Recovery Console, insert the Windows Server 2003 CD-ROM and type ***cd-drive*:\i386\winnt32 /cmdcons** on the command line. The Setup Wizard will install the 8 MB console in a hidden folder called Cmdcons, and will modify the boot.ini file to provide the Recovery Console as a startup option during the boot process.

Removing the Recovery Console

If you ever decide to remove the Recovery Console, you must delete files and folders that are "super hidden." From Windows Explorer, choose the Folder Options command from the Tools menu. Click the View tab, select Show Hidden Files and Folders, clear Hide Protected Operating System Files, click OK and, if you are prompted with a warning about displaying protected system files, click Yes.

Then, delete the Cmdcons folder and the Cmldr file, each of which are located in the root of the system drive. You must next remove the Recovery Console startup option from Boot.ini. Open System from Control Panel, click the Advanced tab, click the Settings button in the Startup And Recovery frame, then, in the Startup And Recovery dialog box, under System startup, select Edit. Boot.ini will display in Notepad. Remove the entry for the Recovery Console, which will look something like this:

```
c:\cmdcons\bootsect.dat="Microsoft Windows Recovery Console" /cmdcons
```

Save the file and close Boot.ini.

Using the Recovery Console

After you have installed the Recovery Console, you can boot the system and select Microsoft Windows Recovery Console from the startup menu. If the console was not installed or cannot be launched successfully, you can restart using the Windows Server 2003 CD-ROM and, at the Welcome To Setup screen, press R to select Repair. The loading takes significantly longer from the CD-ROM, but the resulting Recovery Console is identical to that installed on the local system.

Once the Recovery Console has started, as shown in Figure 13-3, you will be prompted to select the installation of Windows to which you wish to log on. You will then be asked to enter the Administrator password. You must use the password assigned to the local Administrator account, which, on a domain controller, is the password configured on the Directory Services Restore Mode Password page of the Active Directory Installation Wizard.

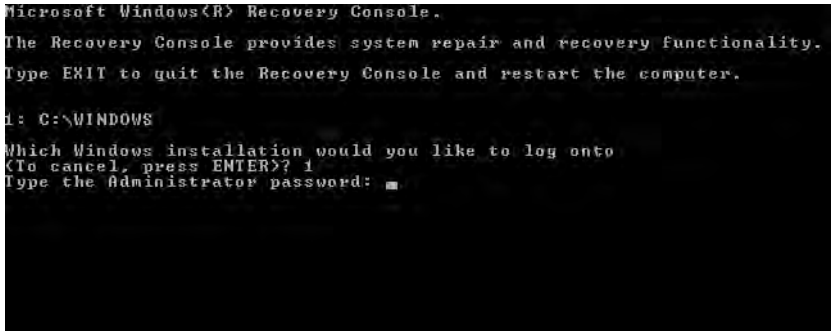

Figure 13-3 The Recovery Console

You can type **help** at the console prompt to list the commands available in the Recovery Console, and **help *command name*** for information about a specific command. Most are familiar commands from the standard command-line environment. Several of the commands deserve particular attention:

- **Listsvc** Displays the services and drivers that are listed in the registry as well as their startup settings. This is a useful way to discover the short name for a service or driver before using the Enable and Disable commands.
- **Enable/Disable** Controls the startup status of a service or driver. If a service or driver is preventing the operating system from starting successfully, use the Recovery Console's Disable command to disable the component, then restart the system and repair or uninstall the component.
- **Diskpart** Provides the opportunity to create and delete partitions using an interface similar to that of the text-based portion of Setup. You can then use the Format command to configure a file system for a partition.
- **Bootcfg** Enables you to manage the startup menu.

The Recovery Console has several limitations imposed for security purposes. These limitations can be modified using a combination of policies (located in the Computer Configuration, Windows Settings, Security Settings, Local Policies, Security Options node of the Local Computer Policy console) and Recovery Console environment variables.

- **Directory access** You can only view files in the root directory, in %*Windir*% and in the \Cmdcons folder. Disable this limitation by setting the policy Allow Floppy Copy And Access To All Drives And All Folders, and using the command **set AllowAllPaths = true**. Be sure to include the space on either side of the equal sign when typing the set command.
- **File copy** You can only copy files to the local hard disk, not from it. Use the policy mentioned above and the command **set AllowRemovableMedia = true**. Be sure to include the space on either side of the equal sign when typing the set command.
- **Wild cards** You cannot use wildcards such as the asterisk to delete files. Implement the policy mentioned above then, in the Recovery Console, type the command **set AllowWildCards = true**. Be sure to include the space on either side of the equal sign when typing the set command.

Practice: Recovering from System Failure

In this practice, you will back up the System State and create an Automated System Recovery Set. You will also install and use the Recovery Console to troubleshoot driver or service failures. Finally, if you have access to a second physical disk drive, you will be able to perform Automated System Recovery to restore a failed server.

Exercise 1: Back Up the System State

1. Log on to Server01 as Administrator.
2. Open the Backup Utility.
3. If the Backup And Restore Wizard appears, click Advanced Mode.
4. Click the Backup tab and select the check box next to System State. Also click the System State label so that you can see the components of the System State listed in the other pane of the dialog box.
5. Type a file name for the backup file, such as **C:\SystemState.bkf**.
6. Start the backup.
7. When the backup is complete, examine the file size of the System State backup file. How big is the file?

Exercise 2: Create an ASR Set

This exercise requires a blank floppy disk and approximately 1.5 GB of free disk space. If you have a second physical disk in Server01, direct the backup to that disk so that you can perform an Automated System Recovery in Exercise 4.

1. Open the Backup Utility. If the Backup And Restore Wizard appears, click Advanced Mode.
2. Click Automated System Recovery Wizard, or choose ASR Wizard from the Tools menu.
3. Follow the prompts. Back up to a file called ASRBackup.bkf on the C drive or, if you have a second physical disk, on that volume.
4. When the backup is complete, examine the file size of ASRBackup.bkf. How big is it? How does its size compare to that of the System State backup?

Exercise 3: Installing and Using the Recovery Console

1. Insert the Windows Server 2003 CD-ROM.
2. Click Start, Run, and then type the following command in the Open box:

   ```
   D:\i386\winnt32.exe /cmdcons
   ```

   where *D:* is the drive letter for your CD-ROM. The Recovery Console will be installed on the local hard disk.
3. To simulate a service in need of troubleshooting, open the Services console from Administrative Tools. Locate the Messenger service. Double-click the service and choose Automatic as the Startup Type.
4. Restart the server.
5. When the server presents the startup boot menu, select Microsoft Windows Recovery Console.
6. When prompted, type **1** to select the installation of Windows Server 2003.
7. Type the password for the local Administrator account.
8. When the Recovery Console prompt appears (by default, C:\Windows>), type **help** to display a list of commands.
9. Type **listsvc** to display a list of services and drivers. Note that the short name of many services is not the same as the long name. However, the short name of the Messenger service is also Messenger. Confirm that its startup is set to Automatic.

10. Type **disable messenger** to disable the service. The output of the command indicates the success of the command and the original startup configuration for the service (in this case, SERVICE_AUTO_START). You should always make note of this setting, so that once troubleshooting has been completed you can return the service to its original state.
11. To quit the Recovery Console, type **exit** and press Enter.

Exercise 4: Restoring a System Using Automated System Recovery

Warning This exercise requires a second physical disk on which an ASR backup has been created in Lesson 2. This exercise will delete all data on the physical disk that contains the system and boot partition. Do not proceed if you have stored any data that you cannot afford to lose.

1. Power off your computer.
2. Restart the computer and open the computer's BIOS. Make sure the system is configured to start from the CD-ROM.
3. Insert the Windows Server 2003 installation CD-ROM.
4. Restart Server01. Watch carefully and, when prompted, press a key to start from the CD-ROM.
5. Early in the text-mode setup phase, setup prompts you to press F2 to run an Automatic System Recovery. Press F2.
6. You will then be prompted to insert the Windows Automated System Recovery disk into the floppy drive. Insert the floppy disk you created in Exercise 2 and press any key to continue.
7. Text-mode setup prepares for Automated System Recovery and a minimal version of the operating system is loaded. This step will take some time to complete.
8. Eventually, a Windows Server 2003 Setup screen will appear.
9. Windows Server 2003 Setup, partitions and formats the disk, copies files, initializes the Windows configuration and then prepares to restart.
10. Remove the floppy disk from the disk drive and allow the computer to restart.

    The installation will continue. When the installation completes, the computer should be restored to its previous state.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. If you are unable to answer a question, review the lesson materials and try the question again. You can find answers to the questions in the "Questions and Answers" section at the end of this chapter.

1. You're setting up a backup job on a computer running Windows Server 2003. You want to back up the registry, startup files, and the COM+ Class Registration database. Which backup option should you select?
   a. %*Windir*%
   b. %*Systemroot*%
   c. System State
   d. None of the above. You cannot back up the registry.
2. You install a scanner on a computer running Windows Server 2003. When you try to restart your computer, the operating system will not start. Which of the following would be the least invasive recovery method to try first to restore the system to operation?
   a. Automated System Recovery
   b. Recovery Console
   c. Safe mode
   d. Directory Services Restore mode
3. A hard disk on a server running Windows Server 2003 has failed. You replace the disk, boot the system, initialize the disk, and create an NTFS volume on the new disk. You now want to restore that data from the last backup job from the old disk. How should you restore the data?
   a. Use the Recovery Console to copy data to the disk.
   b. Use the Backup utility to launch the Restore Wizard.
   c. Use the ASR backup to restore the data.
   d. Use the Last Known Good Configuration option in Safe mode to set up the new disk.

4. A file server on your network will not start. After exhausting all other options, you have decided to use Automated System Recovery (ASR) to recover the system. You created an ASR backup immediately after you installed Microsoft Windows Server 2003 and another one two months ago after you installed a device driver. You perform a full backup of data files once a week. What will ASR restore? (Choose all that apply.)
   a. Data files two months ago
   b. Data files at the last full backup
   c. Disk configuration
   d. Operating system
   e. System State two months ago
   f. System State at the last full backup

Lesson Summary

- The System State includes the registry, startup files, COM+ Class Registration Database, and other service-specific critical system files. It is wise to plan a backup strategy that coordinates backing up the System State along with the system and boot volumes.
- Automated System Recovery uses a setup-like process to return a computer to operation, and then starts a restore operation to recover files from the ASR backup set. It is a recovery process that should be used to restore a system when other less invasive methods, such as Safe mode or the Recovery Console, have been ineffective.
- The Recovery Console is a text-mode command interpreter that allows you to access the hard disk of a computer running Windows Server 2003.

Exam Highlights

Before taking the exam, review the key points and terms that are presented below to help you identify topics you need to review. Return to the lessons for additional practice and review the "Further Readings" sections in Part 2 for pointers to more information about topics covered by the exam objectives.

Key Points

- The System State can be backed up using the Backup Utility or the command prompt, but must be backed up locally. You cannot back up the System State on a remote machine. However, you can back up the local System State to a file on a remote machine, which can then transfer that file to another backup medium.

- To restore the System State on a domain controller, you must restart the domain controller in Directory Services Restore Mode. The System State includes Active Directory. By restoring the domain controller's System State, you are performing a non-authoritative restore, and the domain controller will use standard replication mechanisms to bring itself back up to date. If you want to replicate objects from the restored data to other domain controllers, you must use Ntdsutil to perform an authoritative restore before restarting the domain controller to normal operation.
- Automated System Recovery relies on a catalog of system files stored on the ASR floppy disk to restore files from the Windows Server 2003 CD-ROM, and a comprehensive ASR backup. You prepare the ASR backup set and floppy using the ASR Wizard in the Backup Utility. To perform an Automated System Recovery, restart with the Windows Server 2003 CD and press F2 when prompted.
- The Recovery Console allows you to perform targeted repairs for certain causes of system failure. You can replace system files and disable problematic drivers or services. You can also perform a subset of other system maintenance tasks. The Recovery Console can be launched from the Windows Server 2003 CD or by installing the console on the server's hard drive using the **winnt32 /cmdcons** command.

Key Terms

System State A collection of critical system components including the registry, COM+ Class Registration Database, and startup files. The System State components can be backed up using the Backup Utility or the Ntbackup command. You cannot back up the components separately.

Automated System Recovery (ASR) A new feature that replaces the Emergency Repair process in earlier versions of Windows. Automated System Recovery returns a system to operation by reinstalling the operating system and restoring System State from an ASR backup set.

Recovery Console A utility that provides command-line access to system files and a subset of commands to perform surgical repairs on a failed system.