

Microsoft

Server 2008

Server Administrator (70-646)



**Smarter
Training**

This LearnSmart Exam Manual covers the most important concepts you need to know in order to pass the Windows Server 2008, Server Administrator exam (70-646). By studying this guide, which includes material related to Server 2008 R2, you will become familiar with an array of exam-related content, including:

- Planning for Server Deployment
- Monitoring and Maintaining Servers
- Planning Application and Data Provisioning
- Planning for Business Continuity and High Availability
- And more!

Give yourself the competitive edge necessary to further your career as an IT professional and purchase this exam manual today!

Windows Server 2008 Server Administrator (70-646) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC.
Product ID: 013084
Production Date: September 9, 2011

All rights reserved. No part of this document shall be stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein.

Warning and Disclaimer

Every effort has been made to make this document as complete and as accurate as possible, but no warranty or fitness is implied. The publisher and authors assume no responsibility for errors or omissions. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this document.

LearnSmart Cloud Classroom, LearnSmart Video Training, Printables, Lecture Series, Quiz Me Series, Awdeeo, PrepLogic and other PrepLogic logos are trademarks or registered trademarks of PrepLogic, LLC. All other trademarks not owned by PrepLogic that appear in the software or on the Web Site (s) are the property of their respective owners.

Volume, Corporate, and Educational Sales

Favorable discounts are offered on all products when ordered in quantity. For more information, please contact us directly:

1-800-418-6789
solutions@learnsmartsystems.com

International Contact Information

International: +1 (813) 769-0920

United Kingdom: (0) 20 8816 8036

Table of Contents

<i>Abstract</i>	8
<i>What to Know</i>	8
<i>Tips</i>	8
Planning for Deployment	10
Deploying Windows Server With or Without Hyper-V	10
<i>Server Core Installation</i>	10
<i>Installation Type</i>	10
<i>Upgrade Considerations</i>	11
<i>Preparing for an Upgrade</i>	11
<i>The Installation Process</i>	12
BitLocker Implementation	12
The Initial Configuration	13
Windows Activation Methods	14
Automated Deployment	14
<i>Using the Windows Deployment Services</i>	15
<i>Configuring the Windows Deployment Services</i>	15
Multicast Installation	16
<i>Creating a Multicast Installation</i>	16
DHCP Server Roles	16
<i>Installing the DHCP Server Role</i>	16
The DNS Server Role	17
<i>Global Zones</i>	18
<i>Installing the DNS Server Role</i>	18
<i>Enabling Global Names</i>	18
Active Directory	19
<i>Read Only Domain Controllers</i>	19
<i>Fine Grained Security Policies</i>	19
<i>Architectural Changes</i>	19
<i>Active Directory Functional Levels</i>	19
<i>Installing the Active Directory Domain Services Role</i>	20
<i>Promoting the Server to a Domain Controller</i>	21
The File Server Role	21
<i>The Distributed File System</i>	21

<i>The Windows Search Service</i>	22
<i>Sharing Files</i>	22
<i>The File Server Resource Manager</i>	23
<i>Quotas</i>	23
<i>Creating Quotas</i>	23
<i>File Screening</i>	24
<i>Creating a File Screen</i>	24
<i>Storage Reporting</i>	25
<i>The Windows Search Service</i>	25
<i>Configuring the Windows Search Service</i>	25
DFS Namespaces	26
<i>Implementing a DFS Namespace</i>	27
<i>DFS Replication</i>	28
<i>Replication Topologies</i>	28
Implementing DFS Replication	28
The Print Services Role	30
<i>Adding the Print Services Role</i>	31
<i>Adding a Network Printer</i>	31
<i>Adding a Printer to the Active Directory</i>	32
Application Server and Services.....	32
<i>Adding the Required Roles</i>	32
Server Virtualization	33
<i>Hyper-V Requirements</i>	34
<i>Additional Considerations</i>	34
<i>Physical to Virtual Conversions</i>	35
<i>Virtual Licensing</i>	35
<i>Microsoft's Support Policy</i>	36
<i>Installing Hyper-V</i>	36
<i>Virtual Server Resource Allocation</i>	37
<i>Network Adapters</i>	37
<i>Storage</i>	37
<i>Creating a Virtual Machine</i>	38
<i>Modifying the Virtual Machine's Properties</i>	38
<i>Virtual Hard Disk Types</i>	39

<i>Snap Shots</i>	39
<i>Installing the Integration Services</i>	40
<i>Removable Media</i>	40
The Terminal Services	41
<i>Planning for the Terminal Services</i>	41
<i>Terminal Service Licensing</i>	41
<i>Terminal Server Management</i>	42
<i>Installing the Terminal Services</i>	43
<i>The TS Session Broker</i>	44
<i>Establishing Group Membership</i>	45
<i>Adding Terminal Servers to a Farm</i>	45
<i>Draining a Terminal Server</i>	46
<i>The TS Gateway</i>	46
<i>TS Gateway Policies</i>	46
<i>Configuring the TS Gateway</i>	47
<i>Configuring Client Computers to Use a TS Gateway Server</i>	48
<i>Terminal Services Remote Applications</i>	48
<i>Hosting an Application with TS RemoteApp</i>	49
<i>Making Applications Available to End Users</i>	49
<i>Terminal Service Web Access</i>	50
<i>Configuring TS Web</i>	50
<i>Making a RemoteApp Available Through TS Web</i>	50
<i>Configuring TS Web Access</i>	50
<i>Application Deployment</i>	51
<i>Application Virtualization</i>	51
<i>System Center Configuration Manager</i>	52
Change Management	52
<i>Application Maintenance</i>	53
<i>The Windows Server Update Service</i>	53
Maintaining Windows Server 2008	55
<i>Management Tools</i>	55
ServerManagerCMD	55
<i>Windows PowerShell</i>	56
Remote Management	56

<i>Installing the Remote Server Administration Tools</i>	57
<i>Remote Desktop</i>	57
<i>Enabling Remote Desktop</i>	57
Server Core	57
Server Security	58
<i>The Windows Firewall</i>	58
<i>Auditing</i>	59
<i>Delegation of Administration</i>	59
<i>Active Directory Delegation</i>	60
<i>Performing an Active Directory Delegation</i>	60
<i>Removing a Delegation</i>	61
<i>Server Management Delegation</i>	61
<i>Delegating Server Management</i>	61
Group Policies	62
<i>The Central Store</i>	62
<i>Group Policy Hierarchies</i>	63
<i>Resultant Policies</i>	63
<i>Backing Up Group Policy Objects</i>	64
<i>Restoring a Group Policy Object</i>	64
Monitoring Servers	65
<i>The Event Viewer</i>	65
<i>Monitoring Server Performance</i>	66
<i>The Reliability Monitor</i>	67
<i>The Performance Monitor</i>	67
<i>The Windows System Resource Manager</i>	68
The Certificate Services	69
<i>Types of Certificate Authorities</i>	69
<i>Planning a PKI Hierarchy</i>	70
<i>Installing an Enterprise Certificate Authority</i>	70
Network Access Control	71
<i>Network Access Protection</i>	72
<i>Enabling Network Access Protection</i>	73
<i>Setting Up Network Access Protection Policies</i>	74
<i>Configuring System Health Validators</i>	75

<i>Remote Access</i>	75
<i>Configuring Routing and Remote Access</i>	76
Fault Tolerance	78
<i>Network Load Balancing</i>	78
<i>Failover Clustering</i>	79
<i>The Majority Node Set</i>	79
Backup and Recovery	80
<i>Active Directory Recovery</i>	80
<i>The Active Directory Snapshots and Mounting Tool</i>	80
<i>Recovering Encrypted Data</i>	81
<i>BitLocker Encryption</i>	81
Windows Server Backup	82
<i>Installing Windows Server Backup</i>	82
<i>Creating a Scheduled Backup</i>	83
<i>Restoring Data</i>	83
Troubleshooting	84
<i>The Event Viewer</i>	84
<i>The Performance and Reliability Monitor</i>	84
<i>The Device Manager</i>	84
<i>The Memory Diagnostic Tool</i>	84
<i>Security Troubleshooting</i>	85
Practice Questions	86
Answers & Explanations	96

Abstract

Congratulations on your decision to enhance your career prospects by adding the Microsoft Certified IT Professional (MCTIP): Server Administrator certification to your resume! Once you pass Exam 70-646 (Pro: Windows Server 2008, Server Administrator), you will have earned credit towards this particular certification.

According to Microsoft, becoming certified helps validate real-world skills, which is viewed favorably by many hiring managers. In addition, certified individuals report salary increases and higher productivity. This particular certification helps validate that an individual possesses the necessary skills to serve as database administrator or enterprise messaging administrator.

Although not intended as a replacement for a comprehensive training course, this exam manual provides an overview of all of the major subjects covered on the 70-646 exam. Focusing primarily on material that you must know in order to pass Exam 70-646, the manual was designed with candidates like yourself in mind. To ensure that you get the most out of this exam manual, it has been organized to reflect the actual format of the exam. In other words, the material covered in each chapter reflects topics and objectives tests in a particular domain.

In addition, this guide covers material related to Windows Server 2008 R2. In particular, you will find invaluable information about the R2 update in sections related to upgrade considerations, Active Directory, automated deployment, auditing and storage.

What to Know

Exam 70-646 tests a candidate's ability to administer Windows Server 2008 and Windows Server 2008 R2. When preparing to take the exam, you should make sure that you are familiar with the various server roles and how they are used. You should also make sure that you are well versed in various management related tasks such as patch management, performance monitoring, and system troubleshooting. In addition, you should be familiar with the actual format of the exam. The domains of this particular exam are the following:

- Planning for Server Deployment (19%)
- Planning for Server Management (23%)
- Monitoring and Maintaining Servers (20%)
- Planning Application and Data Provisioning (19%)
- Planning for Business Continuity and High Availability (19%)

Tips

Exam 70-646 tends to be one of Microsoft's easier certification exams; but that doesn't mean that some of the questions may not be tough. There are several things that you can do to improve your odds of passing this exam.

First, don't overthink the questions. Microsoft does not put trick questions on their exams. The correct answer is 100% correct and the other answers are 100% wrong. You will never run into a situation in which two answers are correct but one choice is more right than the other.

Most of the exam questions are multiple choice questions. If you have trouble figuring out an answer, try to eliminate two of the possible choices. That way, if you have to take a guess you can base your guess on the answers that are most likely to be correct.

If you get stuck on an answer, mark it and come back to it later on. You may be able to use some of the other exam questions to figure out the answer that you are stuck on.

If you still can't figure out a question, don't be afraid to guess. If you completely skip a question, that question is scored as if you answered it incorrectly. If you guess at an answer there is at least a chance that you could get the answer right.

Finally, manage your time well. You should have more than enough time to get through the exam; but, don't spend too much time on any one question. If you get stuck, come back to the question at the end of the exam. You want to make sure that you have time to answer all of the questions on the exam.

Planning for Deployment

There are several considerations that must be taken into account when planning for a Windows Server deployment. These considerations are largely based around choosing an edition of Windows Server that will meet your deployment goals.

Deploying Windows Server With or Without Hyper-V

One of the first decisions that must be taken into account is whether or not the server that you are licensing will require Hyper-V. Hyper-V is Microsoft's virtualization solution. It allows a Windows Server machine to host multiple virtual machines. Running Hyper-V requires a 64-bit edition of Windows Server 2008 Standard, Enterprise, or Data center. Hyper-V can be used with either a full installation or a server core installation. If you do not plan on running Hyper-V, then Microsoft offers versions of Windows Server 2008 without Hyper-V for a reduced price. If you plan to run Windows Server 2008 R2, it is available only in 64-bit editions, which all can support Hyper-V.

Server Core Installation

Another decision that you will have to make is whether your deployment will run a full installation or a server core installation of Windows Server. A server core deployment is a type of Windows Server installation in which most of the GUI components are not installed on the server. Instead, the server runs a bare minimal set of services that are designed for performing very specific core tasks. For example, a server core deployment might be used as a DNS server or as a DHCP server.

The primary advantage to a server core deployment is that core servers have a much smaller attack surface; therefore are more secure than a full deployment. Likewise, server core installations often require less memory and CPU time than full server deployments, which makes server core deployments ideal for use in virtual server environments.

However, there are disadvantages. The biggest downside to server core appointments is that core servers must either be managed remotely or by using the command prompt.

Another disadvantage to server core deployments is that server core is primarily designed for infrastructure servers (DNS, DHCP, Active Directory, etc.). As such, most applications will not run on a server core deployment. Server core will not run .NET, which prevents applications such as Exchange Server and SQL Server from running. It also limits what you can do with IIS.

Installation Type

You will also have to decide what type of installation you want to perform. The most trouble-free type of installation tends to be a clean install. A clean install means that you are starting from scratch. You would perform a clean install in situations in which you are setting up a new server or if you have an existing server but want to rebuild Windows from scratch.

Another installation option is a migration. A migration is done when you are going to be replacing an existing server with a new one. A migration starts out with a clean installation of Windows Server, but differs from a clean installation in that once Windows has been installed applications and data are moved or copied from the old server to the new server. After a migration has been performed client computers must be redirected to the new server, unless the old server is taken offline and the new server is given the same name as the old server.

A third option for deploying Windows Server is to perform an upgrade. An upgrade is a process in which Windows Server 2008 is installed on top of an existing Windows Server installation. By upgrading, the new installation will adopt all of the applications and settings that are currently running on the server that is being upgraded.

Upgrade Considerations

Not all legacy Windows Server deployments can be upgraded to Windows Server 2008. There are a number of different factors that determine whether or not a server can be upgraded. One such factor is the operating system that is running on the server. Windows Server 2008 only supports upgrades from Windows Server 2003 SP1 or newer (or from Windows Server 2003 R2). Windows Server 2008 R2 adds an additional limitation in that it can only be upgraded from 64-bit editions of previous Windows Server versions because it only ships as a 64-bit edition. Windows Server 2008 will allow upgrades from the same architecture (for example, 32-bit or 64-bit).

You must also take into account which edition of Windows the server is currently running. Microsoft will allow you to upgrade to either the same edition or a higher edition of the operating system. For example, if the existing servers run Windows Server 2003 Standard Edition, you can upgrade to either Windows Server 2008 Standard Edition or Enterprise Edition. On the other hand, if the server is currently running Windows Server 2003 Enterprise Edition, you can upgrade to Windows Server 2008 Enterprise Edition, but not to Standard Edition.

There are a couple of situations in which Microsoft does not allow an upgrade, even though the versions may match. You cannot upgrade from Windows Server 2003 Web Edition to Windows Server 2008 Web Edition. Likewise, you cannot perform Itanium upgrades, nor can you upgrade to a server core deployment.

As was stated previously, another limitation is that you cannot perform a cross architecture upgrade. Due to this limitation, you cannot upgrade a server that is running a 32-bit operating system to a 64-bit operating system. Likewise, you cannot upgrade a 64-bit operating system to a 32-bit operating system.

Preparing for an Upgrade

Before you can upgrade to Windows Server 2008, you must ensure that your applications are compatible with the new operating system. There are a number of applications that work fine on Windows Server 2003, but that will not run on Windows Server 2008.

Another task that must be done prior to performing an upgrade is to back up the system. In spite of Microsoft's best efforts, upgrades do not always go according to plan; so, it is critical that you ensure that you have a way of rolling back the upgrade if something goes wrong. Keep in mind that simply backing up a server before upgrading it is not enough. You must test your backup to make sure that it can be restored should the need arise. Microsoft does provide a rollback option for failed upgrades, but once you have logged onto the server that was upgraded the rollback option is no longer available. Your only option for rolling back the server at that point is to restore a backup.

You should also make sure to back up any printers that have been defined on your server. The upgrade process removes any printers from the server. In order to be able to print again, the printers must be set up from scratch, or a backup of the printers must be restored.

Depending on the type of upgrade that you are performing, you may need to upgrade your Active Directory schema. You can perform this upgrade with a built-in tool called ADPREP. Running this tool extends the Active Directory schema so that it will support the new Windows Server 2008 features.

The Installation Process

The installation process is quite a bit different from what was found in previous versions of Windows Server. The entire installation process is now graphical. To install Windows Server 2008, follow these steps:

1. Boot the server from the **installation media**.
2. **Setup** will prompt you to specify which language, time zone, and keyboard or input method you want to use.
3. **Click Next**.
4. **Click Install Now**.
5. At this point, you must select the edition of Windows Server 2008 you want to install. You should choose the edition for which you have purchased a license.
6. **Click Next**.
7. **Accept** the license agreement.
8. **Click Next**.
9. You will see a screen asking you whether you want to perform an upgrade or a custom installation. Assuming that you are setting up a new server, choose the **Custom (Advanced) option**. If you need to perform an upgrade you will have to run Setup from within your previous operating system. You cannot perform an upgrade by booting from the DVD.
10. You will now be asked to choose the disk on which you want to install Windows. This screen allows you to create, delete, and format partitions. You can also use this screen to load any drivers that are required in order for Windows to access your RAID array or other storage devices. Select the partition on which you want to install Windows.
11. **Click Next**.
12. The installation process will now complete automatically.

BitLocker Implementation

Unlike the Encrypting File System (EFS), BitLocker can encrypt an entire volume. In order to use BitLocker, your server must be equipped with a Trusted Platform Module (which is also referred to as a TPM chip) or a USB flash drive. The TPM or flash drive stores the encryption keys for the encrypted hard drive and also ensures that the boot environment has not been tampered with. If the boot environment has changed then the TPM chip will prevent the computer from booting from the hard drive. In this instance, you can still boot the server by entering a BitLocker Recovery Key.

Although BitLocker is designed to be used in conjunction with a TPM chip, the TPM chip is not an absolute requirement. A USB flash drive can be used in place of a TPM chip, but the flash drive must be inserted into the server any time that it is booted.

BitLocker is able to encrypt the system drive as well as data drives. To encrypt the system drive, the drive must be equipped with a small partition on which BitLocker related information can be stored. Windows Server 2008 R2 creates this partition automatically if you are performing a clean installation, but Windows Server 2008 does not. However, if you are upgrading from another operating system to Windows Server 2008 R2, this partition will not exist and BitLocker will be unable to encrypt the system drive without extensive adjustments with third-party partition management tools.

If you need to manually prepare a new system for BitLocker encryption, boot the system from the Startup disk, then specify your language, time zone, and keyboard or input method. After doing so, click Repair Your Computer. This causes Windows to boot to a screen showing the operating systems that are currently installed on your system. Since no operating systems exist, just load any necessary RAID drivers and click Next. Finally, click on the Command Prompt option.

WARNING: The method used in this section will erase all information on the hard disk. Use this process with a new system only.

When the Command Prompt window loads, you can partition the system drive for BitLocker by entering the DiskPart command. When you do, the command prompt will change to a DiskPart prompt. You must now select the system drive. To do so, enter the following command:

```
Select Disk 0
```

Now, enter the word Clean. This causes DiskPart to delete any partitions that may exist on the disk. Be careful using the Clean command, as it does not prompt you for confirmation before wiping out every partition on the selected disk. Now, it's time to create the partition that will be used by BitLocker. To do so, enter this command:

```
Create Partition Primary Size=1500
```

The next step in the process is to assign the S drive to the partition that you have just created (this assignment is only temporary). To do so, enter this command:

```
Assign Letter=S
```

Finally, make the partition active by entering the Active command.

Now that the BitLocker partition has been created, you can create the partition to which the operating system will be installed. You can do so by entering this commands:

```
Create Partition Primary  
Assign Letter=C
```

Once your partitions have been created, use the List Volume command to make sure that all of the partitions have been created correctly. Assuming that all is well, go ahead and exit out of DiskPart and format the C: drive and the S: drive. You can accomplish this by entering these commands:

```
Exit  
Format C: /y /q /fs:ntfs  
Format S: /y /q /fs:ntfs
```

The Initial Configuration

After installation is completed, Windows boots and you see a screen telling you that you must change your password. Click okay and you are prompted to enter and confirm a new password. This is the password used by the local administrator account.

After the password is changed, you are signed into Windows Server 2008 as the Administrator. Windows now displays a screen asking you to perform the initial configuration. The screen contains links that you can use to set the server's time zone, configure networking, provide a computer name and domain, enable automatic updates, download updates, and install roles and features. Of course all of these tasks can be performed manually, but the Initial Configuration Tasks screen provides you with a centralized console that you can use to quickly and easily complete all of these tasks.

Windows Activation Methods

After Windows Server 2008 has been installed and configured, the server needs to be activated. Activation is a process that Microsoft uses to ensure that you are not running a counterfeit copy of Windows. If you fail to activate Windows within the allotted amount of time, you will lose server functionality.

The Windows Server 2008 installation process does not ask for a product key. When Windows is installed, it is initially placed into a 60 day evaluation mode. This means that a product key must be provided and Windows must be activated sometime within the next 60 days.

Microsoft provides several methods for activating Windows Server 2008 and Windows Server 2008 R2. One of the most common methods involves using a Multiple Activation Key (MAK). When this method is used you will receive a single product key that can be used to activate multiple computers (up to the number of licenses that you have purchased). During the activation process the server will either communicate directly with Microsoft or with a MAK proxy.

A MAK Proxy is used for activating non Internet connected computers. The MAK proxy connects activation information, sends the information to Microsoft, then relays the activation codes to the servers that need to be activated.

Another method for activating Windows involves using the Key Management Service (KMS). The KMS is a service that runs on a server within your organization and tracks activation requests. Rather than communicating with Microsoft directly, computers requiring activation communicate with the KMS server.

The KMS is designed for use in larger organizations. The service will not begin processing activation requests until it has received requests from either 25 Windows client computers or 5 Windows server computers.

Automated Deployment

Microsoft provides several methods for automatically deploying Windows Server 2008 and Server 2008 R2. In situations where you need to deploy only a few systems, the easiest way to accomplish an automated installation is to use the installation DVD and an answer file. The three methods are discussed below:

1. An answer file is a special extensible markup language (XML) file that provides answers to the questions that Setup asks during the installation process. To perform this type of deployment, copy the answer file to removable media (such as a USB flash drive), insert it into the server, and boot from the Windows installation DVD. Setup will automatically locate and use the answer file.
2. Microsoft also provides a way of installing the answer file to a network share. In this type of deployment, the Windows installation files are copied to a network share along with an answer file (although the answer file can technically reside on removable media too). Rather than booting from the Windows DVD, Setup is run from the network share.
3. The third method of deploying Windows Server 2008 involves using the Windows Deployment Services (WDS). To do so, you will have to download the Windows Automated Installation Kit (WAIK). After doing so you will use a sub-component called the Windows System Image Manager (WinSIM) to create a system image and an answer file. The answer file must be named autounattend.xml and must reside on the media containing the Windows system image. WinSIM can be used to create answer files for any installation type.

Using the Windows Deployment Services

If you plan to use the WDS to deploy Windows, you must have an Active Directory environment and you must be running DHCP and DNS on your network. Another requirement is that the servers you are provisioning must support preboot execution environment (PXE) boot. PXE boot is a capability that is integrated onto the network interface card (NIC). When no installation media is present, the PXE boot capabilities allow the computer to connect to the network and retrieve an IP address from the DHCP server. The system then uses DNS to locate the server that is running WDS.

WDS is actually a server role. The server that is hosting this role must have its disk volumes formatted with NTFS.

To install the WDS, follow these steps:

1. Open the **Server Manager**.
2. Select the **Roles container**.
3. Click **Add Roles**.
4. When the Add Roles Wizard opens, **click Next**.
5. Select the **Windows Deployment Services** check box and **click Next**.
6. **Click Next**.
7. **Accept** the **default role services** and **click Next**.
8. **Click Install**.
9. **Click Close**.

Configuring the Windows Deployment Services

Once WDS has been installed, it must be configured. To do so, select the Windows Deployment Services command from the Administrative Tools menu. When the Windows Deployment Services console opens, follow these steps:

1. Navigate through the console tree to Windows Deployment Services | Servers | <your server>.
2. **Right click** on the listing for your server and choose the Configure Server command from the **shortcut menu**.
3. When the Windows Deployment Services Configuration Wizard opens, **click Next** to bypass the wizard's Welcome screen.
4. The following screen asks where you want to store the image files that you want to deploy to new computers on your network. Provide a path to a local NTFS volume with plenty of free space and **click Next**.
5. If you choose to accept the default path you will see a warning message telling you that you are about to use the system volume. If you see this warning, **click Yes** to continue.
6. The following screen asks you what types of clients you want the server to respond to. The server can respond to pre-staged clients (clients that have been pre-staged in the Active Directory), unknown clients (clients that have not been pre-staged), or both. Make your selection and **click Finish**.
7. The following screen will ask you if you want to add images to the Windows Deployment Server now. Make sure that the check box is selected and **click Finish**.
8. You will now be taken to a screen that asks you for the path to an image file. **Click the Browse button** and browse to your Windows Server 2008 DVD's Sources folder and **click OK**, followed by **Next**.
9. When prompted, provide a name for the image group that you are creating and **click Next**.
10. Verify the summary information and **click Next**. Windows will copy the images to the path that you specified earlier.

Multicast Installation

Multicast installations use the Transport Server component of WDS to multicast images to clients. This deployment method is designed to minimize the amount of network bandwidth that is consumed by the deployment process because the deployment image is streamed to all of the clients at the same time.

Two different types of multicast installations can be configured. The first option is to perform an auto-cast. An auto-cast allows computers to join the multicast installation at any time. If a client joins a multicast that is already in progress, it will accept the remaining portion of the multicast then wait for the multicast to restart so that any missing information can be copied. The second option is a scheduled cast. A scheduled cast allows you to choose when a multicast deployment will begin.

Creating a Multicast Installation

To create a multicast installation, follow these steps:

1. **Open** Server Manager.
2. **Navigate** through the console tree to Roles | Windows Deployment Services | Servers | <your server> | Multicast Transmissions.
3. **Right click** in the console's center pane and choose the Create Multicast Transmission command from the shortcut menu.
4. When prompted, enter a descriptive name for the multicast transmission that you are creating and **click Next**.
5. Select the image group and the deployment image that you want to deploy and **click Next**.
6. Choose whether you want to perform an auto-cast or a scheduled-cast and **click Next**.
7. **Click Finish**.

DHCP Server Roles

The Dynamic Host Configuration Protocol (DHCP) is designed to provide IP addresses and other network configuration settings to network clients. Each DHCP server maintains an IP address scope. A scope is a range of IP addresses that the DHCP server can provide. When a network client requests an IP address from a DHCP server, the DHCP server leases an unused address from its scope to the client. The client does not own this address, but rather borrows it for a limited period of time.

Although DHCP servers provide a critical network service, they were not designed for redundancy or fault tolerance. That being the case, Microsoft recommends that you use what is known as a split scope. A split scope refers to a technique in which the IP addresses that you want to assign are divided into two separate scopes which are leased by two separate DHCP servers. That way, if a DHCP server fails, another DHCP server will still be available to lease IP addresses. If you are running DHCP on virtual servers, it is important to make sure that the two servers making up your split scope are not running on the same virtualization host (physical server). Otherwise, the host can become a single point of failure.

Installing the DHCP Server Role

Before installing the DHCP Server role, you must ensure that at least one of your server's network adapters has been assigned a static IP address. You can install the DHCP Server role on Windows Server 2008 by following these steps:

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** on the Add Roles link.
4. When the Add Roles Wizard starts, **click Next** to bypass the welcome screen.

5. **Select** the check box next to the DHCP Server role and **click Next**.
6. **Click Next**.
7. **Select** the check boxes for the network adapters that you want to use for receiving IP address requests, and **click Next**.
8. On the following screen enter your parent domain name and the IP addresses of your preferred and alternate DNS servers. Specifying an alternate DNS server is optional but recommended.
9. **Click** the Validate buttons to ensure that the DHCP server is able to communicate with the DNS servers.
10. **Click Next**.
11. Tell Windows whether or not WINS is required for applications on the network. WINS is a legacy service that is seldom used with Windows Server 2008, so you will usually choose the WINS is Not Required for Applications on This Network option.
12. **Click Next**.
13. You must now create a scope for the DHCP server. **Click** the Add button. You will be prompted to enter a scope name, a starting address, an ending address, a subnet mask, an optional default gateway, and a subnet type. The subnet type controls the lease duration.
14. Verify that the Activate this Scope check box is selected and **click OK**.
15. If necessary, create any additional scopes. When you are done, **click Next**.
16. Choose whether or not you want to enable DHCPv6 stateless mode for the server. DHCPv6 stateless mode is used for assigning IPv6 addresses.
17. **Click Next**. If you have enabled DHCPv6 stateless mode, you will be taken to a screen which asks you to create an IPv6 scope. Otherwise you are taken to the DHCP Server Authorization screen.
18. Use either your current credentials or supply an alternate set of credentials to authorize the DHCP server in the Active Directory and **click Next**.
19. Take a moment to review the summary screen to ensure that all of the configuration options are correct, then **click Install**.
20. When the installation process completes, **click Close**.

The DNS Server Role

The Domain Name Service (DNS) maps names to IP addresses and IP addresses to names. When a user attempts to access a Web site, they enter the site's URL. This URL is sent to a DNS server which retrieves the requested Web site's IP address.

DNS servers are not just used for Internet access. DNS is a required service for Active Directory environments. In an Active Directory environment, computer names and IP addresses are grouped into zones and are stored in zone files.

One special type of zone that is often covered on the exam is a stub zone. Stub zones are used to keep a parent DNS server up to date when a child domain exists that uses its own DNS servers. Stub zones are only used within your own Active Directory environment.

Another concept that you need to be familiar with is a conditional forwarder. Your DNS server is not aware of every external domain in existence. There are simply too many of them. When someone on your network attempts to access an external domain for which the DNS server does not have a record, the DNS server typically forwards the request to either a higher level DNS server or to a root server. The request moves up the DNS hierarchy until the name resolution request is ultimately resolved.

Needless to say, it can take a long time for name resolution requests to be resolved. If you have domains that you query frequently, you can set up a conditional forwarder that will direct your DNS server directly to the DNS server for the domain that you are querying. This saves time because you do not have to rely on root hints during the name resolution process for that domain.

In Windows Server 2008, DNS zone information is typically stored in the Active Directory. This makes it easy to replicate DNS information to other DNS servers that you may deploy later on. DNS does support zone files instead of Active Directory storage, but this is not a recommended implementation for the DNS zones that provision Active Directory. Only Active Directory integrated DNS implementations provide the security needed by modern Microsoft networks.

Global Zones

Even though DNS is essential for resolving Fully Qualified Domain Names (FQDNs), you may occasionally need to resolve NetBIOS names into IP addresses. In the early days of Windows Server, Microsoft used a service called WINS to resolve NetBIOS names into IP addresses. Although there are some legacy applications that may require WINS, WINS is not required for Windows Server 2008 networks.

Windows 2008 DNS servers allow you to create special zones called Global Zones. Global Zones allow you to map aliases to specific paths on your network. By doing so, users can enter NetBIOS names as a way of accessing portions of your Intranet. For instance, a user might just type the word Marketing or Finance into their Web browser. A record in the Global Zone would translate the request into a valid network path.

Installing the DNS Server Role

To install the DNS Server role, follow these steps:

1. **Open** Server Manager.
2. **Select** the Roles container.
3. **Click** Add Roles.
4. **Click** Next to bypass the Welcome screen.
5. **Select** the DNS Server check box.
6. **Click** Next.
7. **Click** Next.
8. **Click** Install.

Enabling Global Names

If you are going to create a Global Zone, the zone must be Active Directory integrated. Furthermore, all of your DNS servers must be running Windows Server 2008 or higher. It is also worth noting that Microsoft does not support dynamic updates for global names.

Global names are not enabled by default. To enable global names, you must open a Command Prompt window and enter the following command:

```
DNSSCMD <server name> /CONFIG /ENABLEGLOBALNAMESSUPPORT 1
```

Once global names have been enabled, create a DNS zone named GlobalNames. You can then begin creating C Name records in the GlobalNames zone.

Active Directory

Like Windows 2000 Server and Windows Server 2003, Windows Server 2008 supports the use of Active Directory. While many aspects of Active Directory remain unchanged, Microsoft has introduced several new Active Directory features. The method used to set up domain controllers has also changed.

Read Only Domain Controllers

One such change is that it is now possible to setup Windows to act as a read only domain controller (RODC). In Windows NT, only one domain controller (the primary domain controller) was writable. All other domain controllers were read only. In Windows 2000 Server and Windows Server 2003, Microsoft used a multi master domain controller model in which all domain controllers were writable. Windows Server 2008 still adheres to the multi master domain controller model, but also allows the use of RODCs.

RODCs are used primarily in branch office locations. They provide the convenience of an on-premise domain controller, without the security risks of having a fully writable copy of the Active Directory in a location that may not be physically secure.

Fine Grained Security Policies

Another new feature is fine grained security policies. In previous versions of Active Directory, the password policy applied to all of the user accounts in the entire domain. Fine grained security policies make it possible to apply password policies on a per group or a per user basis rather than using a single blanket policy for everyone. These policies are not configured within Group Policy like the traditional password policies. Instead, they are configured using the ADSI Editor as low-level Active Directory objects.

Architectural Changes

Microsoft has also made some architectural changes to the Active Directory. The Active Directory now exists as a service. This means that it is possible to start and stop the Active Directory without having to reboot the domain controller. This is good news for anyone who occasionally needs to take a domain controller offline for maintenance, but doesn't want to have to bring down the entire server.

Microsoft has also improved auditing within the Active Directory. In previous versions of Windows, if an administrator changed some attributes for an Active Directory object, the change could be recorded; but the old values were not retained. Windows Server 2008 allows you to log changes in such a way that the old and new values both appear within the audit log.

Active Directory Functional Levels

With each new version of Windows Server, Microsoft has introduced new Active Directory features. Since older versions of Windows Server do not support the new features, Microsoft has introduced the concept of functional levels as a way of controlling the behavior (and feature availability) of the Active Directory.

There are two different types of functional levels. Domain functional levels control the Active Directory features that are available within an individual domain. There are three domain functional levels available in Windows Server 2008. These domain functional levels include:

1. Windows 2000 Native
2. Windows Server 2003
3. Windows Server 2008

In order to use the new features described in the previous section, the domain functional level must be set to Windows Server 2008. Keep in mind that setting the domain functional level to Windows Server 2008 prevents legacy domain controllers from being added to the domain. Choosing a domain functional level requires careful consideration because you cannot downgrade to a lower functional level once you have enabled a functional level. However, you *can* upgrade.

For example, if you began with a domain functional level of Windows Server 2003, you could not change the domain functional level to Windows 2000 Native, but you could change it to Windows Server 2008. Once it has been set to Windows Server 2008 however, you cannot revert to Windows Server 2003.

The other functional level that you need to be aware of is the forest functional level. The forest functional level dictates the minimum domain functional level. For instance, if the forest functional level is set to Windows Server 2008 then all domains must be configured to use the Windows Server 2008 domain functional level. As a general rule it is best to set the forest functional level to Windows Server 2003 because the Windows Server 2008 forest functional level does not add any additional capabilities beyond what were introduced in Windows Server 2003.

Windows Server 2008 R2 adds an additional functional level. This functional level includes all of the same features as are found in the Windows Server 2003 forest functional level, but also adds the Active Directory Recycle bin, which adds the ability to restore deleted directory objects while the Active Directory is running.

Installing the Active Directory Domain Services Role

In Windows Server 2008, the Active Directory related binaries are installed by adding the Active Directory Domain Services Role. Adding this role does not convert the server to a domain controller, but it does provide the server with the binaries that are necessary for the server to act as a domain controller.

If you plan to promote the server to a domain controller, you must add this role. In previous versions of Windows Server you could promote a server to domain controller status simply by running the DCPROMO command. In Windows Server 2008, DCPROMO is still used, but it will not work until the Active Directory Domain Services Role has been installed.

Even if you do not plan to use the server as a domain controller, adding this role can be helpful because doing so installs the Active Directory management tools.

To install the Active Directory Domain Services Role, follow these steps:

1. **Open** the Server Manager.
2. **Click** on the Roles container.
3. **Click** Add Roles.
4. **Click** Next.
5. **Select** the Active Directory Domain Services check box.
6. **Click** Next.
7. **Click** Next.
8. **Click** Install.
9. When the installation completes, **click Close**.

Promoting the Server to a Domain Controller

After the Active Directory Domain Services role has been installed you can promote the server to domain controller status by following these steps:

1. **Enter** the DCPROMO command at the server's Run prompt.
2. When Windows launches the Active Directory Domain Services Installation Wizard, **click Next**.
3. **Click Next**.
4. **Choose** whether the domain controller will be used within a new forest or an existing forest and **click Next**.
5. If you choose to join the domain controller to an existing forest, you will be prompted to specify whether you want to add the new domain controller to an existing domain or create a new domain.
6. **Click Next**.
7. Type the name of any domain controller in the forest where you plan to install the new domain controller.
8. Tell Windows whether you want to use your current credentials or if you prefer to specify an alternate set of credentials.
9. **Click Next**.
10. You will now be prompted to select the domain that you want to join the domain controller to. Make your selection and **click Next**.
11. You may see a warning message telling you that read only domain controllers cannot yet be installed. This warning asks you if you want to continue. **Click Yes**.
12. Select the Active Directory site in which the new domain controller should be placed and **click Next**.
13. **Choose** the additional options that you want to use for the domain controller. Using check boxes you can designate the domain controller to act as a DNS Server, global catalog server, read only domain controller, or any combination of the three.
14. **Click Next**.
15. Confirm your Active Directory database paths and **click Next**.
16. Enter and confirm your Active Directory Services Restore Mode password and **click Next**.
17. Take a moment to review your configuration settings and **click Next**.
18. **Select** the Reboot on Completion check box.

The File Server Role

The File Server Role gives administrators the ability to centrally manage all of the mechanisms that can be used for sharing files with network clients. Basic file sharing and permissions still work the same as they always have, but the file server role includes more than just basic file sharing.

The Distributed File System

One of the main features that is included in the File Server Role is the Distributed File System (DFS). The DFS provides the ability to create a logical view of all of the files and folders, even if they are scattered across multiple servers. The DFS also allows files and folders to be replicated to other servers as a way of improving efficiency in accessing the files and as a way of providing a degree of fault tolerance.

The Windows Search Service

The File Server Role also includes the Windows Search Service. The Windows Search Service is a replacement for the Indexing Service that was found in Windows Server 2003. As the name implies, this service allows users to perform enterprise wide searches through the Windows Desktop Search feature that is found in Windows Vista and Windows 7.

In the event that you have applications that depend on Windows 2003 services such as the Indexing Service, you can install an optional component called the Windows Server 2003 File Services. This service is not installed by default, and is only intended for backward compatibility purposes.

Sharing Files

Sharing files and folders works the same way in Windows Server 2008 that it did in Windows Server 2003, but Microsoft has redesigned the management tools that are used in the file sharing process. One of the new features in Windows Server 2008 is something called Access Based Enumeration. The Access Based Enumeration feature makes it so that users only see the files and folders that they have access to. If a user does not have access to a file or folder, Windows hides it from view.

Another consideration that must be taken into account with regard to file sharing is whether or not you want to allow offline file caching. Offline file caching is a feature by which folders and their contents can be automatically copied to a mobile user's laptop so that the user will have access to the files regardless of whether or not they are connected to the network. When the mobile user connects to the network then any changes that have been made to the folder's contents are automatically synchronized.

You can share a folder on a file server by following these steps:

1. **Right click** on the folder that you want to share, then choose the Share command from the shortcut menu.
2. Go to the resulting properties sheet's Sharing tab and click the **Advanced Sharing button**.
3. The share name is automatically populated, but you can specify an alternate share name if necessary by using the **Add button**.
4. Specify the number of simultaneous connections that will be allowed for the share.
5. Enter a comment regarding what the share will be used for.
6. **Click** the Permissions button.
7. Specify the permissions that should apply to the share and **click OK**.
8. **Click** the Caching button.
9. Set the offline caching options for the shared folder. You can disable caching, set the caching to automatic, or set the folder so that only the files that are manually specified by the user are cached. This is the default behavior.
10. **Click OK** three times.

It is worth noting that file shares can also be managed directly through the Server Manager. To do so, follow these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | File Services | Share and Storage Management.
3. Upon selecting the Share and Storage Management container, Windows will display all of the shares that exist on the server (including the default system shares).
4. Use the Provision Share link to create a share.

The File Server Resource Manager

The File Server Resource Manager is a utility for managing the quantity and type of data stored on Windows file servers. The File Server Resource Manager offers three main features including:

- **Quotas** – Quotas can be used to control how much data a user can store. Quotas can be applied at the volume level or at the folder level.
- **File Screens** – File Screens allow you to control the types of files that are allowed to reside within a share. For example, a file screen could be used to prevent users from storing MP3 files within a network share.
- **Reporting** – Although quotas and file screens are often thought of as mechanisms for restricting users, they can be used as a means for gathering statistical information without actually restricting anything.

Quotas

Storage quotas in Windows Server 2008 are based on the use of templates. Templates define quota limits and the action that will occur when the quota threshold is exceeded. There are two types of quotas – hard quotas and soft quotas.

A hard quota restricts storage. If a hard quota is exceeded it can either prevent the user from storing any additional data or it can prevent anyone from storing any additional data within a volume or folder.

Soft quotas are purely informational in nature and provide a means for telling users that they are consuming too much disk space, but without placing restrictions on the user.

When you define a quota, you can link an event to that quota. An event is an action that occurs when the quota is exceeded. For example, you could configure Windows to automatically launch a disk clean up script or to send someone an E-mail message when a quota has been exceeded.

Windows Server 2008 also supports a special type of quota called an auto quota. An auto quota is applied to a parent folder and then automatically creates a quota on any sub folders that are created beneath that folder. For example, suppose that you wanted to create a quota on each of your User folders. Rather than manually applying a quota for each user, you could set an auto quota on the Users folder then a quota would automatically be created for all of the sub folders. This includes currently existing and future sub folders.

The reason why it is possible to create auto folders is because of the way that Windows Server 2008 quotas are created. Quotas are based on templates. If you make a change to a template then all quotas that were based on that template are automatically updated to reflect your change.

Creating Quotas

You can create a quota by following these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | File Services | Share and Storage Management | File Server Resource Manager | Quota Manager | Quotas.
3. **Click** Create Quota.
4. When the Create Quota dialog box opens, enter the quota path.
5. Tell Windows whether you want to create a quota directly on the path or apply an auto quota.
6. **Choose** the quota template that you want to use or define a set of custom quota properties.
7. **Click** the Create button.

File Screening

File screening is a technology by which it is possible for network administrators to block certain types of files. When you configure Windows to block a file type, you must generally apply the block to a file group, which is a collection of file extensions that are used for similar purposes. For instance the file group related to audio files might include MP3 and WMA files.

In addition to using file groups, it is also possible to specify your own naming patterns. For instance, if you wanted to block every file beginning with the letter A, you could block A*.*.

File screening also supports the use of exclusion patterns. For example, suppose that you wanted to block all .AVI files except for some Windows Server 2008 training videos, all of which start with the letters WS2008. You could accomplish this by blocking *.AVI and creating an exclusion for WS2008*.AVI.

One important thing to remember about file screening is that the screen is based solely on the file's name. Windows does not attempt to look inside the files to verify the type of file. Therefore, if you were to create a filter to block MP3 files, a user could bypass the filter by renaming all of their MP3 files to use a different file extension.

File screening supports active and passive blocking. Active blocking is a hard block that prevents users from saving certain types of files. A passive block won't stop a user from saving a file, but it will trigger an action of your choosing. For example, you might create a passive file block that sends an E-mail to the HR department if someone saves an MP3 file to the network.

Building a file screen involves combining some templates. The File Groups template specifies the type of file that is to be blocked. The Screening Type template lets you choose between active and passive blocking, and the Actions template lets you control what happens when a blocked file type is detected. When you create a file screen, that file screen applies to the selected folder and to all sub folders. However, it is possible to apply exceptions at the sub folder level if necessary.

Creating a File Screen

You can create a file screen by following these steps:

1. **Open** Server Manager.
2. Navigate through the console tree to Server Manager | Roles | File Services | Share and Storage Management | File Server Resource Manager | File Screening Management | File Groups.
3. Check to see if any of the default file groups will meet your needs. If not, you can edit any of the default file groups or create your own file groups.
4. After reviewing the file groups, select the **File Screen Templates** container.
5. Check to see if any of the existing file screen templates will meet your needs. You will most likely have to create your own file screen template by clicking the **Create File Screen Template link**.
6. When the **File Screen Template Properties** sheet opens, enter a name for the template that you want to create.
7. **Choose** whether you want to perform active screening or passive screening.
8. **Choose** the file groups that you want to block.
9. **Click** OK.
10. **Select** the File Screens container.
11. **Click** the Create File Screen link.
12. When the **Create File Screen dialog box** opens, enter the file screen path.
13. **Choose** the file screen template on which you want to base the file screen.
14. **Click** the Create button.

It is worth noting that when you create a file screen, it will only detect new instances of blocked files. If files of blocked file types already exist within the specified path then those files will not be affected by the file screen.

Storage Reporting

If you have implemented quotas and file screens then it is a good idea to use storage reporting to ensure that the mechanisms you have implemented are working correctly. You can also use storage reporting to identify storage trends or to be alerted of various issues (such as low disk space). Some of the other things that can be reported on include duplicate files or files that are accessed most frequently.

Creating a storage report typically requires setting up a report task. A report task contains the report settings (the data that you want to include in the report) as well as the schedule for the report. Although not encouraged by Microsoft, it is possible to run a one-off report without creating a report task.

The Windows Search Service

The Windows Search Service is a File Server role feature that is designed to take the place of the Indexing service. However, the Windows Indexing Service still exists and can be installed if necessary. The Windows Indexing Service should be used only for backward compatibility with applications that depend on it. Unless such applications exist, file servers should run the Windows Search Service instead. The Windows Search Service and the Indexing Service cannot both be run on the same server.

When run on a server, the Windows Search service indexes the files that are stored on the server. The desktop version of the service also indexes non-file data such as E-mail messages. Once the server's contents have been indexed, Windows Vista and Windows 7 clients will be able to use Windows Desktop Search to quickly search the server's contents and Windows 7 machines may include the networked resources in a local Library. Windows Desktop Search can also be used on machines running Windows XP or Windows Server 2003, but Windows Desktop Search is an optional component on these operating systems and must be manually installed.

The Windows Search Service can be configured to index either individual folders or entire volumes. Microsoft recommends that you avoid volume level indexing unless the volume is used exclusively for storing data.

Configuring the Windows Search Service

In order to configure the indexing locations, follow these steps:

1. **Open** the Control Panel.
2. **Double click** on the Indexing Options icon.
3. When the **Indexing Options dialog box** opens, click the **Modify button**.
4. **Select** the folders and / or volumes that you want to index.
5. **Click** OK.

The Indexing Options dialog box also contains an Advanced button. If you click this button, Windows will open the Advanced Options dialog box. This dialog box will allow you to perform some basic configuration tasks such as choosing whether or not you want to index encrypted files or whether or not you want to treat similar words as different words.

The Advanced Options dialog box also contains mechanisms for rebuilding your index, moving the index location, and selecting the types of files that you want to index.

DFS Namespaces

In larger organizations, it is rare to be able to store all of a user's files on a single server. Typically, such organizations may have many different file shares spread across a lot of different file servers. Although necessary, this type of file location structure can be confusing for end users because they may have trouble remembering where to find various types of files.

This is where the Distributed File System (DFS) comes into play. The DFS allows shares that are spread across multiple file servers to be grouped into a single logical file system. That way users can easily locate files on the network without having to worry about where the files are physically stored. These logical structures are known as DFS namespaces.

There are two different types of DFS namespaces that an administrator can create. These are known as:

- **Standalone Namespace** – A standalone namespace uses the DFS server as a root. Network clients would access the namespace by going to \\<server name>\namespace. The advantage to creating a standalone namespace is that standalone namespaces are the only type of DFS namespaces that can be used in conjunction with failover clustering. Therefore, if you need to make a file server's contents highly available then a clustered standalone DFS namespace is the way to go.

It is worth noting that Windows Server 2008 Standard Edition can only host a single standalone namespace per server. Windows Server 2008 Enterprise Edition and Datacenter Editions can host an unlimited number of standalone namespaces on each server.

- **Domain-Based Namespace** – A domain-based namespace uses the domain name as the namespace root. For example, a domain-based namespace might be named \\contoso.com\namespace.

In order to understand the benefits of domain-based namespaces, you need to understand the concept of DFS targets. A DFS namespace typically points to a target, which serves as the root folder for the DFS namespace. With that in mind, domain based namespaces offer a tremendous benefit when it comes to migrating data from an old file server to a new one.

In the past, migrating data to a new file server meant that any links to the file share had to be updated to reflect the new location. This could be a tremendous undertaking. With domain based namespaces, however, this is no longer necessary. All of the existing connections to the share will continue to function (assuming that the connections are linked to the DFS namespace). The administrator can simply edit the DFS root to point to a target on the new file server rather than having to manually redirect each user.

Another advantage to domain based DFS namespaces is that they can provide redundancy. It is possible to set up multiple DFS servers and tie them all into the Active Directory so that they receive information about the DFS structure that has been established. For example, a company that has several branch offices might place a DFS server in each branch office. When a user attempts to access a file, the Active Directory uses site information to determine which DFS server is in the closest physical proximity to the user. The user's request is then directed to that server.

If one of the DFS servers should fail, users are automatically redirected to another DFS server. Technically, redundant domain-based DFS servers are different from failover clusters; but, they offer many of the same benefits.

In order to achieve the maximum benefit from domain-based namespaces, the domain that is hosting the DFS servers must be set to Windows Server 2008 functional level or higher. Furthermore, all of the DFS servers must be running Windows Server 2008 or later. Once those requirements have been met, the DFS namespace can be converted from Windows 2000 mode to Windows Server 2008 mode. Windows 2000 mode offers basic domain-based DFS functionality and it allows you to mix and match DFS server types, but it does not offer all of the benefits that you can get when you switch to Windows Server 2008 mode.

One such feature that is only found in Windows Server 2008 mode is access based enumeration. As you will recall, access based enumeration makes it so that users can only see the resources that they have been given access to.

The other benefit to using Windows Server 2008 mode is scalability. In the past, organizations that needed to achieve scalability often used standalone namespaces because they could scale to over 5000 folders. This simply wasn't possible with domain based namespaces. With Windows Server 2008 mode however, domain based namespaces offer the same level of scalability as what was previously found only in standalone namespaces.

Implementing a DFS Namespace

Assuming that the File Services role is already installed on your server, a DFS namespace can be implemented by following these steps:

1. **Open** the Server Manager.
2. **Select** the File Services container.
3. **Click** the Add Role Services link.
4. When prompted, select the Distributed File System check box. Both of the sub components (DFS Namespaces and DFS replication) will be selected automatically.
5. **Click** Next.
6. **Select** the option to create a namespace now.
7. Enter a name for the namespace that you want to create.
8. **Click** Next.
9. **Select** whether you want to create a domain based namespace or a standalone namespace.
10. Assuming that you are creating a domain based name space, you will have the option of enabling Windows Server 2008 mode. If your domain and your DFS server do not meet the Windows Server 2008 mode requirements, this option will be grayed out.
11. **Click** Next.
12. The following screen allows you to add folders and folder targets. **Click** the Add button.
13. **Click** Browse.
14. **Select** the folder that you want to use as a target.
15. **Click** OK.
16. **Click** Next.
17. Take a moment to verify that the information presented on the summary screen is correct.
18. **Click** Install.
19. When the installation completes, **click Close**.

DFS Replication

DFS replication allows administrators to synchronize folders across multiple servers. DFS replication is based on a multi-master model. In other words, if a user makes a change to a file, that change can be written to any of the DFS servers and will then be replicated to all of the other DFS servers within the DFS replication group.

One of the nice things about DFS replication is that you can use it to replicate folders that are not even shared. For example, if you have some backend folders that you want to replicate to other servers you can accomplish the task with DFS replication. This further illustrates the point that DFS namespaces and DFS shares are independent of each other. Just because you create a DFS namespace, it does not mean that the namespace must be replicated. Likewise, you can replicate folders regardless of whether or not those folders are a part of a DFS namespace. In fact, Microsoft even uses the DFS replication engine to replicate the SYSVOL folder that is used by the Active Directory.

DFS replication uses a mechanism called Remote Differential Compression (which is also known as delta changes). In other words, when a change is made to a file, the DFS replication service does not replicate the entire file. Instead, only the parts of the file that have been changed are replicated. This helps to conserve network bandwidth. Additionally, it is important that the anti-virus software be compatible with DFS replication or it can cause problems. You should check with the software vendor to verify compatibility.

Replication Topologies

When you create a DFS replication group, you will be asked to specify a replication topology that the group members can use. You are given either two or three choices, depending on the number of replica servers in the group that you are creating. Your choices include:

- **Hub and Spoke** – The Hub and Spoke replication topology option is only available in replication groups consisting of three or more members. In this topology a server is designated as the hub server. All of the other replicas (known as spoke servers) are connected to the hub server and communicate with it directly. Spoke servers communicate only with the hub server, not with other spoke servers. This topology works well in environments in which content is created centrally and needs to be pushed out to other DFS servers on the network.
- **Full Mesh** – A full mesh topology is a topology in which every DFS server within the replica group contains a logical connection to every other server in the group. Microsoft recommends using this topology for replica groups with ten or fewer members. Otherwise the replication process can get bogged down as a result of the excessive numbers of connectors between replication group members.
- **No Topology** – The third option that is available to you is to create a replication group with no topology. You would only use this option in situations in which you wanted to manually create a custom replication topology after the replication group has been created.

Implementing DFS Replication

Assuming that a DFS namespace and at least one target already exist, you can enable DFS replication by following these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | File Services | DFS Management | Namespaces | <your namespace> | <your target>.
3. **Select** the Replication tab.
4. **Click** the Replicate Folder Wizard link.

Another way in which you can set up DFS replication is to follow these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager Roles | File Services | DFS Management | Replication.
3. **Click** the New Replication Group link.
4. When the New Replication Group Wizard starts, select the type of replication group that you want to create. You can create a multipurpose replication group or a replication group for data collection. A multipurpose replication group is just a standard replication group. A replication group for data collection copies data from various servers to a central hub server so that the data can be backed up.
5. After making your choice, **click Next**.
6. Enter a name for the replication group that you are creating.
7. Enter a description of the replication group's purpose.
8. **Verify** that the correct domain is selected.
9. **Click** Next.
10. **Click** the Add button.
11. Specify the names of the servers that you want to include in the replication group. These servers will be known as the replication group's members.
12. If prompted, supply a set of administrative credentials.
13. **Click** Next.
14. **Choose** the replication topology that you want to use.
15. **Click** Next.
16. **Choose** whether you want the replication group members to continuously replicate data with one another or if you prefer for the group members to follow a custom replication schedule.
17. **Choose** whether you want the replication process to take full advantage of your network bandwidth or if you need to conserve bandwidth through the use of throttling.
18. **Click** Next.
19. **Select** one of your DFS servers to act as the replication group's primary member. The primary member is the server that contains the data that you want to replicate. If a replicated folder already exists on multiple servers, the primary member will be authoritative over that folder.
20. **Click** Next.
21. **Click** the Add button.
22. Specify the local path to the folder (on the primary member server) that you want to replicate.
23. Provide a name that will represent the replicated folder.
24. Set any necessary permissions.
25. **Click** OK.
26. **Click** Next.
27. **Click** the Edit button.
28. Set the membership status to **Enabled**.
29. Provide the local path of the folder.
30. **Click** OK.
31. **Click** Next.
32. **Verify** the settings that are displayed on the summary screen.
33. **Click** the Create button.
34. When the replication group has been created, **click Close**.

The Print Services Role

The Print Services Role allows you to manage your organization's printers through a special Windows Server management console.

Keep in mind that if you simply want to print from a server or allow others to print from a printer that is attached to a server, it is not necessary to install the Print Services Role. Basic print functionality is a core operating system feature. The Print Services Role is aimed at managing printers. Installing the Print Services Role also adds support for the Line Printer Daemon (LDP) used for UNIX / Linux printing and the ability to facilitate Internet printing.

As you prepare for the 70-646 exam, it is important to remember the terms that Microsoft uses for printing. These terms come into play on almost all of the printing related exam questions. The terms that you need to know include:

- **Print Device** - The Print Device refers to the physical hardware printer.
- **Printer** - In Microsoft speak, a printer does not refer to a hardware device; rather, to the print driver that gets installed onto the server. Windows views the printer as the logical representation of the print device.

The reason why Microsoft takes this approach is because it allows for some very flexible configurations. For example, it is possible to map multiple printers to a common print device. That way, you can assign a different printer to each group of users and assign a printing priority based on the group's needs. Print jobs received by higher priority printers will print first.

You can also create a print pool by linking multiple print devices to a single printer. This is useful in organizations that do a lot of printing, because multiple print devices are available to juggle the printing workload.

When it comes to facilitating printing for network clients, you generally have two choices. The clients can either communicate directly with the printer, or they can relay print jobs through a Windows Server 2008 server that is running the Print Services role.

It is usually more advantageous to connect network clients to a server that is configured with the print services role. There are several advantages to doing so. First, the Print Services Role makes it possible to publish printers in the Active Directory. This means that users are able to search for a printer that meets their needs.

Another benefit to using the Print Services role is that doing so allows printers to be audited. If for example, someone is burning through excessive amounts of ink and paper you can use auditing to determine who has been doing the most printing.

Yet another benefit to using the Print Services Role is that you can display centralized status information for all of the printers throughout the entire enterprise. It is even possible to scan subnets to locate printers for inclusion in the management process.

Finally, the Print Services role allows for centralized printer management. If for instance you need to install a new printer driver, you can install the driver directly on the server rather than having to deploy it to each network workstation.

Adding the Print Services Role

To install the Print Services Role, follow these steps:

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** the Add Roles link.
4. **Click** Next.
5. **Select** the Print Services check box.
6. **Click** Next.
7. **Click** Next.
8. **Verify** that the Print Server check box is selected. You can also select the LDP Service and / or the Internet Printing check box if you need UNIX / Linux or Internet printing capabilities.
9. **Click** Next.
10. **Click** Install.
11. **Click** Close.

Adding a Network Printer

Once the Print Services Role has been installed you can add a shared printer by completing these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Print Services.
3. **Click** the Add and Share a Printer on the Network link.
4. When the Network Printer Installation Wizard begins, choose the installation method that you want to use. Your options include: Search the Network for Printers, Add a TCP/IP or a Web Services Printer by IP Address or Host Name, Add a New Printer Using an Existing Port, or Create a New Port and Add a New Printer. In most cases, you will use the Web Services Printer by IP Address or Host Name option.
5. **Click** Next.
6. Set the Type of Device option to AutoDetect.
7. Enter the name or IP address for the printer.
8. Make sure that the Auto Detect the Print Driver to Use check box is selected.
9. **Click** Next.
10. Once the printer has been detected, Windows will ask you if you want to use an existing printer driver or install a new driver. Make your selection and **click Next**.
11. Assuming that you have chosen the option to install a new print driver, you must now either select your print driver from the list or click the Have Disk button and supply Windows with the path to the print driver file.
12. **Click** Next.
13. Enter a name for the printer.
14. **Select** the Share This Printer check box.
15. Provide a share name for the printer.
16. Enter the print device's physical location.
17. Provide any necessary comments regarding the print device's capabilities or intended use.
18. **Click** Next.
19. **Click** Next.
20. **Click** Finish.

Adding a Printer to the Active Directory

By adding a network printer to the Active Directory, you make it possible for users to search for the printer. For example, a user might do a search for a printer on the third floor that can handle double sided print jobs in color.

When you set up a network printer, Windows does not automatically add the printer to the Active Directory. If you want to add a printer to the Active Directory, you can do so by following these steps:

1. **Choose** the Print Management command from the server's Administrative Tools menu.
2. Navigate through the console tree to Print Management | Custom Filters | All Printers.
3. **Right click** on the printer that you want to add to the Active Directory and select the Properties command from the resulting shortcut menu.
4. When the printer's properties sheet appears, select the Sharing tab.
5. **Select** the List in Directory check box.
6. **Click** OK.

Application Server and Services

An application server is any server that hosts an application used by your users. Before you can deploy an application server, you need to understand the requirements of the application that the server is going to be running. Typically, an application server will run a line of business application that tends to be demanding from a resource consumption prospective. As such, application servers are often dedicated machines that host the application and nothing else.

Adding the Required Roles

Most applications that are designed to run on top of Windows Server 2008 depend on various operating level system services. As such, you may have to deploy roles or role services before you can install a line of business application.

Although every application has different requirements you will often find that you are required to deploy the Application Server role on your server before you can set out to install a line of business application. To install the Application Server role, follow these steps:

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** the Add Roles link.
4. When the Add Roles wizard begins, **click Next** to bypass the Welcome screen.
5. **Select** the Application Server check box.
6. When prompted, click the **Add Required Features** button.
7. **Click** Next.
8. **Select** any additional role services that you wish to install.
9. **Click** Next.
10. Review the confirmation screen to make sure that everything is correct.
11. **Click** Install.
12. **Click** Close.

There are a lot of different components to the Application Server role; but, the core features that must be installed in order to facilitate the role are version 3.0 of the .NET Framework and the Windows Process Activation Service, which includes the Process Model, the .NET environment, and the configuration APIs. With Windows Server 2008 R2, this changes to version 3.5 of the .NET framework and many applications also require service pack 1 for the 3.5 version.

It is worth noting that although Windows Server 2008 includes version 3.0 of the .NET Framework, this is not the latest version. Although installing the Application Server role requires installing version 3.0 of the .NET Framework you will often find that you are required to install either a service pack or a later version of the .NET Framework prior to installing your server application. This is true when installing SharePoint Server 2010 as an example.

Server Virtualization

One of the biggest trends in IT today is server virtualization. The basic idea behind server virtualization is that server hardware is becoming so powerful that much of a server's CPU, memory, and disk resources are unused. In some cases, it may be possible to make better use of the server's available resources by installing multiple applications or services onto a single server. For instance, a server might be configured to act as a domain controller, a file server, a DNS server, and a DHCP server all at once.

Unfortunately, there are several problems with running multiple applications or services on a single server. One issue that you might encounter is compatibility between services or applications. In some cases two applications might require two different versions of the same DLL file, which would prevent the two applications from running on the same machine.

Even if you don't run into compatibility problems, security can be an issue. The more services and applications that you have running on a single server, the greater the potential attack surface of that server. In addition, some applications may require you to open firewall ports for grant additional rights to a service account. If such actions are performed then it could lead to other applications that are running on the server becoming vulnerable to attack because those other applications are running under the assumption that they are running in a secure environment in which no additional firewall ports of been opened and no additional rights have been granted.

Yet another reason for avoiding running multiple services and applications on a single server is that doing so can greatly complicate the troubleshooting process whenever problems do occur. Often times an application can have adverse effects on another application. Usually these effects are felt right away, but sometimes the problems don't show up until a lot later. In those types of situations the difficulty of troubleshooting the problem can be greatly compounded by the simple fact that multiple services and applications are running.

Finally, as the number of services and applications that are running on a server increases, the server also becomes increasingly more mission-critical. Think about what would happen if a domain controller were to fail. As long as there were other domain controllers that were still running, the failure probably wouldn't be catastrophic. Now imagine the same situation, only this time the domain controller is also serving as a DNS server, a DHCP server, and a file server. If that server were to go down you would probably have a major problem on your hands.

Because of these and numerous other issues, administrators have long preferred to use a dedicated server for each service or application. The problem with doing so is that server hardware can be really expensive. Even if the cost of the server itself is not an issue, physical space within the data center comes at a premium. Eventually the data center is going to fill to capacity and there will not be room for any additional servers.

One great solution to all of these problems is to use server virtualization. Server virtualization works by allowing multiple copies of the Windows operating system to run on a single server. Each one of these operating systems runs within a fully self-contained virtual machine that is oblivious to the other virtual machines. Each virtual machine assumes that it is running on physical hardware.

The Microsoft solution for virtualizing servers involves using a Windows Server 2008 component called Hyper-V. Hyper-V is a hypervisor for Windows Server 2008. What this means is that a small block of code, known as a hypervisor, runs beneath the operating system. Virtual machines communicate with the hypervisor, which in turn communicates with the server hardware. This approach offers far better performance than the emulation model that was used with earlier server virtualization solutions such as Microsoft's Virtual Server.

One of the nice things about Hyper-V is that it greatly simplifies the process of migrating to a new server. Virtual machines are self-contained and portable. Therefore, if you need to replace an aging physical server with a new one, all you have to do is to export the virtual machines, then import them into the new server. You don't have to worry about most of the challenges that normally go along with migrating to new server hardware.

Hyper-V Requirements

There are several requirements that must be satisfied in order to run Hyper-V. First, you will need a 64-bit version of Windows Server 2008 or you will need Windows Server 2008 R2, which comes only in a 64-bit distribution. The Standard, Enterprise, and Data center editions of Windows Server 2008 all include Hyper-V. Hyper-V is not included with the web edition of Windows Server 2008.

Hyper-V also requires a server with hardware level virtualization support. Specifically, this means that the server must support either Intel VT or AMD-V. Data Execution Prevention must also be enabled. Depending on the hardware platform that is being used this means that either the Intel XD bit or the AMD NX bit must be enabled.

Additional Considerations

One of the biggest drawbacks to server virtualization is that by consolidating multiple servers to a single physical machine, you can create a single point of failure. If the physical server were to drop off-line then all of the virtual servers that it was hosting will also fail. When you consider that a single physical server can potentially host dozens of virtual machines you can begin to see how the failure of a single physical machine could result in a major outage.

That being the case, it is important to create your virtual server infrastructure in a way that avoids a single point of failure. In the case of Hyper-V, a best practice is to take advantage of failover clustering. That way, if a host server were to fail, another server within the failover cluster could take over for the failed machine and the virtual servers will remain online and available.

When you build a failover cluster for Hyper-V you must use shared storage. Shared storage is a centralized disk resource which contains all of the virtual hard drive files used by the virtual machines. The host servers within a failover cluster typically connect to a shared storage pool by using iSCSI or Fibre Channel.

Even though a failover cluster is designed to prevent your virtual machines from dropping offline as a result of a hardware failure there are situations which can cause an entire failover cluster to fail. For example, if a power failure in part of the building were to cause multiple cluster nodes to fail then the remaining cluster nodes may not be able to retain quorum (more on that later). This would cause the cluster to fail.

Because failover clusters are not entirely bullet proof it is important to arrange your virtual servers in a way that compromises your planned redundancy. For example, it would be a bad idea to place all of your domain controllers onto a single host server because if that host server failed (or if the cluster failed) then all of your domain controllers would go down with it. It would be better to stagger your domain controllers across multiple hosts.

This brings up another point. If you are planning on virtualizing all of your domain controllers then you should make your host servers part of a workgroup – not part of the domain. Otherwise there are situations that could result in a catch 22 in which no domain controller is able to authenticate the Hyper-V server, but you can't boot the domain controller because it is on the Hyper-V server. If you really need for your host servers to be domain members then you should keep at least a two physical domain controllers on your network. Remember this fact on exam day.

Physical to Virtual Conversions

Although it is easy to think of server virtualization as a solution designed purely for new servers, server virtualization is also excellent for consolidating existing servers. This is especially true for servers that are running on legacy hardware. As hardware ages it becomes increasingly more difficult to find the parts that are necessary to keep it running. Moving a legacy server to virtual hardware is a great way of getting around the challenges of maintaining legacy hardware.

Microsoft refers to the practice of migrating from physical to virtual hardware as a physical to virtual migration (also known as a P2V Migration). There are a number of ways in which a physical to virtual migration can be accomplished. The easiest way is to use standard imaging tools to duplicate an existing server and move its contents to virtual hardware. In addition, Microsoft also offers several solutions for performing physical to virtual migrations. For example, Microsoft's System Center Virtual Machine Manager offers physical to virtual migration capabilities as does the Virtual Server Migration Toolkit.

Virtual Licensing

The subject of licensing virtual servers can be a little bit tricky, and it is something that you should expect see on the exam. As mentioned earlier, 64-bit editions of Windows Server 2008 Standard, Enterprise, Data center all include Hyper-V.

The license for Windows Server 2008 Standard Edition allows you to run one virtual machine that is running Windows Server 2008 Standard Edition without having to purchase any additional software licenses.

Technically, this means that you can install two separate copies of Windows Server 2008 Standard Edition (one on the host machine and another on the guest machine). However, Microsoft's licensing policy prevents the copy of Windows Server 2008 Standard Edition that is running on the host machine from being used for anything other than running Hyper-V. In other words, you cannot use this instance of Windows to run applications or host non-Hyper-V related system services.

The Enterprise Edition of Windows Server 2008 is licensed a bit differently. Each server running Windows Server 2008 Enterprise Edition is allowed to run up to four different virtual machines (each running Windows Server 2008) without incurring any additional licensing costs. Unlike the Standard Edition, the Enterprise Edition license does allow you to use the host operating system for functions other than running Hyper-V. However, Microsoft generally recommends that the host operating system only be used for Hyper-V.

Regardless of whether you are running the Standard Edition or the Enterprise Edition, you are allowed to host more virtual machines than those that are included with the license. You must simply pay to license any additional virtual machines that you choose to run. For example, if you needed to run 10 virtual machines on top of Windows Server 2008 Enterprise Edition then you would need to purchase six licenses. One of those licenses would take care of the host operating system and the first four virtual machines. The other six licenses would be used for the remaining virtual machines.

Microsoft also allows you to run operating systems other than Windows Server 2008 within virtual machines. However, you are responsible for licensing those operating systems.

If you're planning on virtualizing a large number of servers, your best option might be to purchase a copy of Windows Server 2008 Data center Edition. Windows Server 2008 Data center Edition allows you to run an unlimited number of copies of Windows Server 2008 in virtual machines that are hosted on that server. You are also allowed to host non-Windows and legacy Windows operating systems; but, you must license any virtual machines that are not running Windows Server 2008.

Microsoft's Support Policy

With a few exceptions, all of Microsoft's usual support policies apply regardless of whether a server is running on physical or virtual hardware. For example, Windows NT Server has not been supported by Microsoft in years. It is possible to perform a physical to virtual conversion on your Windows NT Servers to allow them to function in a Hyper-V environment, but doing so does not mean that Microsoft will suddenly begin supporting Windows NT just because it is running within a Hyper-V virtual machine.

You should also check the support policies for your applications prior to virtualizing an application server. Some applications (or application configurations) are not supported in virtual server environments. For example, until recently Microsoft would allow you to run Exchange Server 2010 in a virtual server environment, but they would not allow you to virtualize Exchange 2010 servers that were hosting the Unified Messaging role. As such, it is important to check with application vendors to make sure that there are no restrictions or guidelines for running the application in a virtual server environment.

Installing Hyper-V

Before you install Hyper-V you should perform a system update make sure that you have all of the latest patches. While it is always a good idea to make sure that Windows Server 2008 is up to date, it is especially important to do so before installing the Hyper-V role.

The reason for this is that when Microsoft shipped Windows Server 2008, Hyper-V hadn't been officially released yet. As such, the operating system came with a release candidate (beta) version of Hyper-V. Performing an update before you install Hyper-V ensures that you receive the latest version and do not end up installing the release candidate version.

You can install Hyper-V by following these steps:

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** the Add Roles link.
4. When the Add Roles Wizard begins, **click Next**.
5. When Windows asks you what roles you want to install, select the **Hyper-V** check box.
6. **Click** Next.
7. **Click** Next to clear the introduction to Hyper-V.
8. **Select** the network adapters that you would like to bind to virtual networks. Microsoft generally recommends reserving one physical network adapter for the host operating system.
9. **Click** Next.
10. **Click** Install.
11. **Click** Close.
12. **Reboot** the Server.

Virtual Server Resource Allocation

Hosting virtual servers with Hyper-V is all about resource allocation. You have to be able to share your server's physical hardware resources among all of the virtual machines (plus the host operating system) in a way that ensures that each virtual machine receives the resources that it needs without depriving any of the other virtual machines or the host operating system of necessary resources.

When you create or configure a virtual machine in Hyper-V, there are a number of different resources that you can allocate.

Memory – the most important resource to manage in a Hyper-V environment is memory. Your server's physical memory has to be shared among all of the virtual machines, while still saving some memory for the host operating system (which is sometimes called the parent partition). In Hyper-V, memory is allocated, not shared. This means that each virtual machine reserves a block of memory and that the block of memory becomes exclusive to the virtual machine to which it was assigned for as long as the virtual machine is running. Virtual machines do not hold onto memory when they are powered off.

CPU – another resource that you must allocate among your virtual machines is the CPU. Unlike memory, CPU cores can be simultaneously shared by multiple virtual machines. Hyper-V uses a concept known as logical processors. Each logical processor corresponds to a single physical processor core. Therefore a server with twin quad core processors contains eight logical processors. Although it is best to avoid over committing logical processors, Hyper-V will not stop you from allocating more CPU resources than what your server actually has. When you do, Hyper-V will attempt to manage processor affinity in a way that ensures the best possible performance.

Network Adapters

When you install Hyper-V, you will be asked which network adapters you want to use with Hyper-V. When a network adapter has been committed for use with Hyper-V, that adapter is no longer directly accessible to the host operating system. Instead, Hyper-V creates a virtual network switch and links the network adapter to it.

If you look at a network adapter that has been committed for use with Hyper-V from within the host partition, you will see that the only component that is bound to the network adapter is the Microsoft Virtual Network Switch protocol. However, Hyper-V will create a virtual network adapter within the host operating system that is connected to the virtual switch and ultimately to the physical network adapter. Settings such as IP addresses must be assigned to the virtual network adapter.

As a best practice, it is a good idea to reserve at least one network adapter for the host operating system and not commit that adapter for use with Hyper-V. If your host server contains multiple network adapters than you might even consider dedicating a separate network adapter to each virtual server.

Storage

Hyper-V makes use of virtual hard disk files (.VHD files). These virtual hard disk files can be stored either on direct attached storage or on a SAN. It is also possible to use a technique called Pass Through as a way of configuring a virtual machine to use a physical SCSI drive rather than the virtual hard drives which are normally used.

Windows Server 2008 R2 uses VHD files in a new way. You can actually boot a physical machine from a VHD placed on the system. This allows you to multi-boot between several installations on the system using multiple VHD files stored on the physical drive and it can be very useful for testing or training environments.

Creating a Virtual Machine

Hyper-V makes it very easy to create virtual machines. Virtual machines are created and managed through the Hyper-V Manager. Microsoft also makes an enterprise class virtual machine management utility called System Center Virtual Machine Manager, but this tool is not included with Windows Server 2008.

You can create a new virtual machine by following these steps:

1. **Open** the Hyper-V Manager.
2. Navigate through the console tree to Hyper-V Manager | <your server>.
3. **Right click** on the container for your server and select the New | Virtual Machine commands from the shortcut menus.
4. When the New Virtual Machine Wizard launches, **click Next** to bypass the Welcome screen.
5. Enter a name for the virtual machine that you want to create.
6. **Select** the Store the Virtual Machine in a Different Location check box.
7. Supply a path within which you want to store the new virtual machine's configuration files and virtual hard drive files.
8. **Click Next**.
9. Specify the amount of memory that you want to allocate to the new virtual machine.
10. **Click Next**.
11. Specify the network adapter that you want to use with the virtual machine. Only network adapters that have been provisioned for use with Hyper-V are displayed. If you need to use multiple network adapters with the virtual machine, you can add additional network adapters by editing the virtual machine's settings after the virtual machine has been created.
12. **Click Next**.
13. **Choose** the type of virtual hard disk that you want to create. You can create a new virtual hard drive, use an existing virtual hard drive, or attach a virtual hard disk later. The option to attach a virtual hard disk later is useful if you want to create a fixed length virtual hard disk file or if you want to use SCSI pass through.
14. If you are creating a new virtual hard disk, verify the path and then specify the size of the virtual hard disk that you want to create.
15. **Click Next**.
16. **Choose** how you want to install an operating system. You can choose to manually install an operating system later on, or you can install an operating system from a boot CD / DVD. There is also an option to install an operating system from a boot floppy. This can be useful if you are trying to deploy a legacy operating system such as Windows NT Server.
17. **Click Next**.
18. Verify the summary information for the new virtual machine.
19. **Click Finish** to create the new virtual machine.

Modifying the Virtual Machine's Properties

The New Virtual Machine Wizard presents you with a limited number of options for creating a new virtual machine. If you need to perform additional configuration on the virtual machine that you have just created (or if you need to edit another virtual machine) you can do so by right clicking on the listing for the virtual machine and choosing the Settings command from the shortcut menu.

The Settings page provides a number of useful options. For starters, you can use this screen to allocate additional hardware to the virtual machine. For example, if you wanted to provide the virtual machine with multiple virtual network adapters or multiple virtual hard disk files this is where you would do it. You can also use the Settings page 2 to change the number of virtual processors or the amount of memory that is assigned to the virtual machine.

Virtual Hard Disk Types

When you create a new virtual machine, the Hyper-V Manager gives you the option of creating a virtual hard disk, using an existing virtual hard disk, or waiting until later to create a virtual hard disk. Part of the reason for the option to wait to create a virtual hard disk is because Hyper-V actually allows you to create three different types of virtual hard disks.

If you allow the New Virtual Machine Wizard to create a virtual hard disk for you it will create a dynamically expanding virtual hard disk. This means that no matter how large you tell Hyper-V to make the virtual hard disk, the actual size of the .VHD file will initially be very small. As you add data to the virtual hard disk, the file size will grow to accommodate the data until the .VHD file eventually reaches the maximum size that you have specified.

Another option is to create a fixed size virtual hard disk. A fixed size virtual hard disk reserves the disk space that you have specified as the virtual hard disks size. From a performance standpoint, creating a fixed size virtual hard disk is your best bet. There is such a large performance difference that Microsoft requires Exchange 2010 mailbox servers running on Hyper-V to use fixed size virtual hard disks.

Still another advantage to fixed size virtual hard disks is that by creating hard disks of a fixed size you avoid the situation in which you accidentally overcommit your disk resources. If you only use dynamically expanding virtual hard disk files that it is possible to create virtual hard disks that exceed your physical storage capacity. This will work fine for a while; but, as you add data to the virtual hard disks you will eventually run the server out of storage space.

With all of these advantages you might be wondering why Microsoft did not design Hyper-V to create fixed size virtual hard disks by default. The most likely reason for this is that it can take quite a while to create a fixed size virtual hard disk; whereas dynamically expanding virtual hard disks are created almost instantly.

One last type of virtual hard disk that is supported by Hyper-V is a differencing disk. A differencing disk is a virtual hard disk file that has a parent-child relationship with another disk that you want to leave intact. It allows you to make changes to the data were to the operating system without affect the parent disk. That way, you can easily revert the system to the way it existed prior to those changes being implemented.

Snap Shots

One of the main features you need to know about Hyper-V is the snapshot feature. Snapshots are essentially a point in time backup of a virtual machine. When you create a snapshot what you are actually doing is creating a differencing disk. From that point on, any changes that you make to the virtual machine are written to the differencing disk instead of to the virtual machines main virtual hard disk file. That way, if you ever need to roll the changes back you can.

As handy as snapshots can be, there are two important things that you need to know about them. First, snapshots impact performance. When the server needs to perform a read operation it must first check the differencing disk to find out if the data exists there. If the data does not exist on the differencing disk, it is read from the main virtual hard disk file. This can result in a huge performance deficit. The problem is further compounded as additional snapshots are created. Thankfully, snapshots are designed to be temporary and the snapshot data can either be deleted or merged into the main virtual hard disk file.

The other thing the need to know about snapshots is that even though they can be very handy for backing up a virtual machine prior to making a configuration change, they are not suitable for use on database servers. Snapshots can wreak havoc on relational databases. In fact, Microsoft's support policy for Exchange Server 2010 mailbox servers explicitly prohibits the use of snapshots.

Installing the Integration Services

When you install an operating system on a virtual machine, the operating system may lack the drivers necessary to communicate with the virtual hardware. Operating systems running in virtual environments need device drivers to be able to communicate with things like network cards, just as they would if they were running on physical hardware.

Virtual machine device drivers are added to the virtual machines through the use of the Integration Services. The Integration Services can be installed on most newer Windows Server and Windows desktop operating systems. It is worth noting however that the Integration Services are not necessary for virtual machines running Windows Server 2008 because the drivers are already built into the operating system.

Non-Windows operating systems such as Linux or UNIX will not work with the integration services. When such operating systems are in use, Hyper-V uses emulated hardware. Hardware emulation is not as efficient as using the integration services, but it does allow an otherwise incompatible operating system to function from within Hyper-V by emulating a legacy network adapter (the Intel 21140).

You can install the integration services onto a virtual machine by following these steps:

1. **Open** the Hyper-V Manager.
2. If the virtual machine is not already running, **right click** on it and choose the Start command from the shortcut menu.
3. **Double click** on the virtual machine to connect to it.
4. When the virtual machine's operating system finishes installing, log in.
5. **Choose** the Insert Integration Services Setup Disk from the virtual machine's Action menu.
6. Follow the prompts to install the integration services. The exact prompts vary somewhat depending on the operating system that is being used on the virtual machine, but the installation process is extremely simple.

Removable Media

Virtual machines running on top of Hyper-V do support the use of removable media such as CDs or DVDs, but do not support the use of USB devices such as flash drives or external hard drives.

If you want to use a CD or DVD with a virtual machine, follow these steps:

1. **Open** the Hyper-V Manager.
2. **Double click** on the virtual machine within which you want to use the CD or DVD.
3. **Choose** the DVD Drive | Capture commands from the virtual machine's Media menu.

If you receive an error message during this process it usually means that another virtual machine has already captured the DVD drive. You can release the drive by connecting to the virtual machine that is using the drive and selecting the DVD Drive | Uncapture commands from that virtual machine's Media menu.

The Terminal Services

One of the big problems that has plagued network administrators for decades is that of desktop application management. Managing desktop applications can be a major undertaking. Each application has to be deployed in a uniform manner and patches must typically be installed on a regular basis.

In the past, desktop application management was often performed manually. Someone from the helpdesk had to physically travel to each individual PC and perform the necessary software installation or maintenance task. Larger networks have traditionally gotten around that problem by using application management software. Such software packages are able to push applications and application patches to desktop machines, but such tools are often expensive and complicated to use.

One way of addressing the challenges of desktop application management is to take advantage of the Windows Terminal Services. The terminal services allow applications that would normally be installed on a desktop computer to be installed on the server instead. That way, the application can be managed centrally and the administrator must only maintain a single copy of the application, rather than a copy for each individual desktop.

When applications are run in a Terminal Server environment, the application actually runs on the server as opposed to the desktop. Users connect to the terminal server using a special Remote Desktop client. User input such as keystrokes and mouse movements are sent to the terminal server. The terminal server processes those inputs and refreshes the users display.

There are several different advantages to using a terminal server environment. First, as previously mentioned, installing applications on a terminal server greatly simplifies application management.

Another advantage is that the terminal services allow users to run applications that might not ordinarily run on their desktops. In fact, some organizations even go so far as to replace user's PCs with dumb terminals as a way of reducing hardware costs.

Finally, using the terminal services can improve security. If all of the applications are running on a server and users are accessing those applications through a dumb terminal, then it becomes far more difficult (if not impossible) for users to install unauthorized software or copy data from the network.

Planning for the Terminal Services

There are several considerations that must be taken into account for organizations that are planning to use the terminal services. One such consideration is how the Terminal Services will be used. It is only necessary to deploy the Terminal Services (and purchase Terminal Service Client Access Licenses) if you plan to allow end user connectivity to the Terminal Server. If you only plan to use the Terminal Services for remote administration and management tasks, you do not need to install the Terminal Services. You can use Remote Desktop instead.

Terminal Service Licensing

You must also plan for software licensing. As is the case with any Windows Server 2008 deployment, each client will require a Client Access License in order to connect to the server. In addition, however, clients will also need a separate Client Access License to connect to the Terminal Services.

Application licensing can also be a little bit tricky in a Terminal Services environment. As a general rule of thumb, you need an application license for every user who will be accessing the application. However, some applications based their licensing count on the number of computers on which the application is installed rather than on the number of users that will be using the application. From a compliance standpoint, such applications should still be licensed as if they were being installed on individual PCs rather than on a single server. It is also worth noting that there are a few applications with licenses that specifically prohibit Terminal Server installations.

In an effort to ensure that your Terminal Services deployment complies with Microsoft licensing requirements, Microsoft requires you to deploy a Terminal Services License Server. The server's job is to keep track of the number of Client Access Licenses that you have purchased as well as the number of users who are using those licenses. Licensing can be performed on a per-domain or per forest basis.

If you are planning to upgrade from an earlier version of the Terminal Services, the first server you will need to upgrade is the license server. The reason for this is that a Windows Server 2008 Terminal Services License Server is able to manage licenses for earlier versions of the Terminal Services. However, a legacy Terminal Services Licensing Server (such as the one included with Windows Server 2003) is not able to manage licenses for terminal servers that are running Windows Server 2008.

Another important thing to know about licensing the Windows Terminal Services is that licenses can be assigned to either users or devices, but the licensing process is semi-permanent. When a user or device connects to a Terminal Server, the licensing server assigns a Client Access License to that user or device. This license becomes permanently associated with the device or the user to whom it was assigned. In extreme situations there is a way of revoking a license, but Microsoft does not allow you to freely assign and then remove licenses on an as needed basis.

Another area in which planning is essential is security. One thing you must bear in mind about the Windows Terminal Services is that users have to be assigned sufficient rights in order to be able to run the applications. In some cases (especially with older applications) it may be necessary to give users full blown administrative rights just to get the applications to run correctly.

With that in mind, you should never run the Windows Terminal Services on a domain controller. If you give users administrative rights on a domain controller then you are essentially making them domain administrators. That would give the users full blown administrative control over the entire Active Directory. As such, you should always run the Terminal Services on a dedicated server.

Another thing that you need to know about security is that in Windows Server 2008 Microsoft has changed how the logon process works for the Windows Terminal Services. Previously, a user establishes a connection to a Terminal Server using the Remote Desktop Protocol (RDP). Once the connection was established, a desktop session would be created and the user would be presented with a logon prompt.

In Windows Server 2008 Microsoft uses network level authentication to validate the user's computer before a Terminal Server session can even be created. Doing so improves security and it helps to reduce resource consumption on the terminal server. However, support for Network Level Authentication requires that desktop clients be using RDP version 6.0. This version is found in Windows XP SP2 and higher (as well as Windows Vista and Windows 7). RDP clients running version 5.2 and lower will not be able to establish a session with a Windows Server 2008 Terminal Server so long as network level authentication is enabled.

Terminal Server Management

Another important consideration is terminal server management. If an organization is only using a single Terminal Server then the terminal server can be configured manually directly through the server console.

Often, larger environments will use multiple Terminal Servers. In these types of environments it is still possible to configure the servers manually; but, manual configurations make it difficult to ensure a consistent configuration from one server to the next. In these situations it is generally advisable to manage Terminal Server configurations through the use of group policy settings.

Installing the Terminal Services

The first thing you need to know about installing the Windows Terminal Services is that you have to install the Terminal Services before you install any applications. The reason for this is that many applications are not designed to be run in a multiuser environment, which is what the Terminal Services are. As such, the Terminal Services monitor application installation and may take corrective action that allows an application that was designed for a single user environment to run properly within a Terminal Services environment. That noted, not all applications can be made to run within a Terminal Services environment.

Like most other Windows Server 2008 components, the Terminal Services are role based. Before you install the Terminal Services you must decide which roles you want to use. The available roles include:

- **Terminal Server** – The Terminal Server role includes all of the core Terminal Services components.
- **TS Licensing** – Installing the TS Licensing role configures the server to act as a Terminal Service Licensing Server. You must have at least one Terminal Services Licensing Server in your organization if you plan to use the Terminal Services.
- **TS Session Broker** – The TS Session Broker role is a role that allows users to reconnect to a previously used Terminal Service session in a load balanced environment. The session broker accomplishes this by using a table that matches user names to session IDs and the names of the servers on which each session is running.
- **TS Gateway** – The TS Gateway provides RDP tunneling over HTTPS sessions so that remote users can connect to Terminal Service sessions over the Internet.
- **TS Web Access** – The TS Web Access component allows users to use a Web browser to connect to a Terminal Server session or to a RemoteApp application. The Web browser only facilitates the connectivity process. Terminal Service applications are not run through the Web browser, but rather through the Remote Desktop client.

Assuming that you are installing all of the roles, you can install the Terminal Services by following the steps below. Some of these steps may be omitted for deployments involving fewer Terminal Service roles.

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** the Add Roles link.
4. When the Add Roles Wizard's Welcome screen is displayed, **click Next**.
5. When the roles selection screen is displayed, select the **Terminal Services** option.
6. **Click** Next.
7. **Click** Next to clear the introduction to Terminal Services screen.
8. **Choose** the Terminal Service roles that you want to install.
9. **Click** Next.
10. You should now see a warning message telling you not to install applications prior to installing the Terminal Services. **Click Next** to clear this warning.
11. Specify whether or not you want to use Network Level Authentication. You should always enable Network Level Authentication unless you have legacy Terminal Services clients.
12. **Click** Next.
13. **Choose** the licensing model that you want to use. You can assign licenses based on device or user, or you can specify a licensing model later on. It is usually best to license connections on a per device basis because device licenses are not affected by employee turnover. If you choose to configure the licensing method later then you have 120 days before you are required to license your terminal server.

14. **Click** Next.
15. You are now prompted to specify which users or groups will have access to the Terminal Server. The Administrators group is given access by default, but it will be necessary to grant access to anyone else who will require Terminal Server access. Microsoft recommends granting access to groups rather than to individual users.
16. **Click** Next.
17. Specify whether the Terminal Service Licensing Server will provide licenses for terminal servers in the current domain only or for the entire forest.
18. **Click** Next.
19. Provide Windows with an SSL certificate that it can use for encryption.
20. **Click** Next.
21. If you are installing the TS Gateway role, you must then choose to define an authorization policy either now or later.
22. **Click** Next.
23. **Click** Next to clear the Network Policy and Access Services introductory screen.
24. The following screen asks which Network Policy Server role features should be installed. **Click Next** to accept the default selection.
25. **Click** Next to clear the Internet Information Services (IIS) introductory screen.
26. **Click** Next to accept the default IIS role services.
27. Take a moment to review the summary of the configuration information that you have provided.
28. **Click** Install.
29. **Click** Close.
30. **Reboot** the server.

The TS Session Broker

The TS Session Broker, which is one of the Terminal Service roles, performs two major tasks. First, it is responsible for establishing load balancing within the Terminal Services environment. Larger organizations often use multiple Terminal Servers so that no one single server has to support the full workload. In such environments it is the TS Session Broker that is responsible for distributing user sessions across all of the available Terminal Servers in an even manner.

Also, TS Session Broker maintains a list of user sessions. If a user gets disconnected from their session they will be able to reconnect to the same session so they can pick up where they left off. In previous versions of Windows Terminal Services, a disconnected user might reconnect to an entirely different Terminal Server session and possibly on a different server. This would cause the user to lose whatever they were working on at the time that they were disconnected.

When building a Terminal Server farm, you must implement a load balancing solution. Typically, this is performed with either the Network Load Balancing service or by using DNS round Robin. This load balancing solution evenly distributes inbound connection requests among the Terminal Servers in the farm.

Once a user has connected to a Terminal Server, the Terminal Server contacts the TS Session Broker to determine whether or not the user had previously been connected to a session. In that type of situation, the user's connection is redirected to the appropriate server and the user's connection is reestablished.

Even if the user did not previously have a Terminal Service session, the TS Session Broker might still tell the Terminal Server to move the user's connection to a different server within the farm in an effort to maintain an even balance of user sessions across the various servers within the farm.

Believe it or not, the TS Session Broker role does not have to run directly on a Terminal Server. In fact, Microsoft actually recommends that you run the TS Session Broker role on a file server or something similar. That way, the TS Session Broker is not consuming resources on the actual Terminal Server.

Another thing that you need to know about the TS Session Broker is that there are a few requirements that must be met in order for the TS Session Broker to do its job. First, all of the Terminal Servers within the farm must be running Windows Server 2008. Likewise, all of the network clients must support RDP 5.2 or higher. This means that clients can run Windows XP, Vista, or 7.

Establishing Group Membership

The first step that you should take in configuring the TS Session Broker is to make sure that all of the Terminal Servers within the farm belong to the session directory local group. To do so, follow these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Configuration | Local Users and Groups | Groups.
3. **Double click** on the Session Directory Computers local group.
4. When the group's dialog box opens, click the **Add button**.
5. Enter the names of the servers that you want to add to the group.
6. **Click** OK.
7. **Click** OK.

Adding Terminal Servers to a Farm

Once the TS Session Broker is in place, you can create a farm and begin adding Terminal Servers to it. To do so, you can use the steps listed below. The farm is created when you provide a name for the farm (see step 7). You must make sure to use exactly the same farm name for each server or you will end up creating multiple farms.

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Terminal Services | Terminal Services Configuration.
3. In the lower, middle pane **double click** the Member of Farm in TS Session Broker link.
4. When the Properties dialog box opens, go to the **TS Session Broker** tab.
5. **Select** the Join a Farm in TS Session Broker check box.
6. Enter the IP address of the server that is running the TS Session Broker service.
7. Enter the name of the farm that you are joining. The farm name must be identical for every server in the farm.
8. **Select** the Participate in Session Broker Load Balancing check box.
9. Provide a relative weight for the server within the farm.
10. Make sure that the Use IP Address Redirection check box is selected.
11. **Click** OK.
12. **Repeat** the process for the other Terminal Servers in your farm.

Draining a Terminal Server

One of the challenges with working with Terminal Servers is that you can't perform maintenance tasks, such as installing patches, while users are connected to the Terminal Server. When you need to perform maintenance it is necessary to remove user connections from the server by using a process called draining.

Draining a Terminal Server does not actually disconnect any of the user sessions. Instead, it prevents users from reconnecting to the server once they have ended their session. That way, the terminal server will eventually be free from user connections.

To drain a Terminal Server, follow these steps:

1. **Open** the Server Manager on the server that you want to drain.
2. Navigate through the console tree to Server Manager | Roles | Terminal Services | Terminal Services Configuration.
3. **Click** the Member of Farm in TS Session Broker link.
4. When the Properties dialog box appears, select the **General** Tab.
5. **Choose** the appropriate User Logon Mode.
6. **Click** OK.

There are three different User Logon Modes that you can specify in the procedure above. These include:

- **Allow All Connections** – Users are allowed to connect to the server and establish Terminal Service sessions.
- **Allow Reconnections But Prevent New Logons** – Users are allowed to reconnect to existing sessions, but no new sessions will be created.
- **Allow Reconnections But Prevent New Logins Until The Server is Restarted** – This option does exactly the same thing as the Allow Reconnections But Prevent New Logins option except that once the server is rebooted the User Logon Mode will be automatically set back to Allow All Connections.

The TS Gateway

The TS Gateway acts as the gateway into your network for external clients. Clients residing outside of your network are able to tunnel into the network by establishing an RDP over HTTPS session with the TS Gateway server. Windows uses Internet Information Server (IIS) to facilitate the underlying connectivity.

The TS Gateway role is supported on Windows Server 2008 Standard, Enterprise, and Datacenter editions. However, Standard Edition only supports 250 sessions. If you need to allow additional sessions, you will either need to deploy multiple TS Gateway Servers or you will need to use Enterprise Edition.

TS Gateway Policies

The Remote Desktop Protocol (RDP) that is used by the Terminal Services is the same protocol used for managing servers through Remote Desktop. As such, it is imperative that the TS Gateway Server be configured so that only authorized users are allowed to establish an RDP session. Microsoft provides TS Gateway security through the use of policies. There are two policies that you need to be aware of:

- **CAP** – Connection Authorization Policy
- **RAP** – Resource Authentication Policy

As the name implies, the Connection Authorization Policy controls who is allowed to connect to your network through the TS Gateway Server. This policy can be configured to use Network Access Protection (NAP) for even better security.

The Resource Authentication Policy controls which network resources users are allowed to access once they have connected to the TS Gateway Server. This policy provides granular control that you can use to restrict users to accessing only their own desktop.

Configuring the TS Gateway

Before users will be able to connect through the TS Gateway Server, you will have to create a Connection Authorization Policy and a Resource Authorization Policy.

As discussed previously, the Connection Authorization Policy controls connectivity, while the Resource Authorization Policy provides access to network resources. It is a common practice to create a single Connection Authorization Policy and multiple Resource Authorization Policies. For example, a Connection Authorization Policy could be configured to give everyone in the Domain Users group access to the network. You could then use a series of Resource Authorization Policies to control which specific network resources various groups of users are allowed to access.

You can create the necessary policies by completing these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Terminal Services | TS Gateway Manager.
3. **Double click** on the listing for your TS Gateway Server.
4. **Click** on the Connection Authorization Policy link.
5. When Windows displays the New TS CAP dialog box, provide a name for the policy that you are creating.
6. **Select** the properties sheet's Requirements tab.
7. **Choose** the types of authentication methods that you want to allow (Passwords and / or smart cards).
8. **Click** the Add Group button.
9. Supply the name of the group that you want to allow to connect to the TS Gateway Server.
10. **Click** OK.
11. **Select** the properties sheet's Device Redirection tab.
12. **Choose** which devices you want to redirect.
13. **Click** OK.
14. When you are returned to the main Server Manager screen, navigate through the console tree to Server Manager | Roles | Terminal Services | TS Gateway Manager | <your TS Gateway server>.
15. **Click** on the Create Resource Authorization Policy link.
16. When the New TS RAP properties sheet appears, you must provide Windows with a name for the new Resource Authorization Policy that you are creating.
17. **Select** the properties sheet's User Groups tab.
18. **Click** the Add button.
19. Enter the name of the group for which you wish to provide access to network resources.
20. **Click** OK.

21. **Select** the properties sheet's Computer Group tab.
22. Specify which computers or network devices the users that were previously specified should be allowed to access.
23. **Select** the properties sheet's Allowed Ports tab.
24. Verify that the option to allow connections only through TCP port 3389 is selected (unless you need to use a different port for some reason).
25. **Click** OK.

Configuring Client Computers to Use a TS Gateway Server

The TS Gateway Server is only used by clients who are attempting to connect to the Windows Terminal Services from outside of your network. Clients won't be able to simply open the Remote Desktop Client and provide the name of the Terminal Server as they would if they were using a computer on the local network because they are connecting over the Internet. Instead, there is a little bit of configuration that needs to take place on the client computer in order to allow the client to connect through the TS Gateway Server. You can configure the client computer by following these steps:

1. **Open** the Remote Desktop Client on the client computer.
2. Enter the name of the computer that the user wants to connect to. This is done in exactly the same way that it would be if the user resided on the local network.
3. **Click** the Options button.
4. When the Remote Desktop Options are displayed, click on the **Advanced** tab.
5. **Click** the Settings button located in the Connect from Anywhere section.
6. **Select** the Use These TS Gateway Server Settings option.
7. Enter the fully qualified domain name or the IP address of the TS Gateway server.
8. Set the Logon Method option to **Ask For Password** (NTLM).
9. Make sure that the Bypass TS Gateway Server for Local Addresses option is selected.
10. **Click** OK.

Terminal Services Remote Applications

Historically, the Windows Terminal Services have been used to provide users with a full-blown hosted Remote Desktop session. In some cases, however, you might prefer to allow the users to continue to use their own Windows desktop. After all, users can sometimes find it confusing to have to log into Windows then establish a second Windows desktop session once they connect to the Terminal Services.

In Windows Server 2008, Microsoft introduced a new Terminal Services component called TS RemoteApp. This new component makes it possible to use the Terminal Server to host individual applications rather than entire Windows desktops. That way, administrators can still enjoy all of the benefits of centralized application management; but can do so in a way that allows the users to continue to use their own Windows desktops. Remote applications are presented to the end user in a manner that is almost completely transparent. A user would be hard pressed to tell that a hosted application was not installed directly on their desktop.

Like a normal Terminal Server session, the TS RemoteApp component makes use of RDP. The difference is that TS RemoteApp establishes a single application RDP session. Because the Terminal Server is still involved in the application hosting process, the users will still need the appropriate rights for connecting to the Terminal Server.

Hosting an Application with TS RemoteApp

Before you can use the Windows Terminal Services to host individual applications, you must tell Windows that you want to make specific applications available through TS RemoteApp. You can do so by following these steps:

1. **Install** the application onto your Terminal Server if it is not already installed.
2. **Open** the Server Manager on your Terminal Server.
3. Navigate through the console tree to Server Manager | Roles | Terminal Services | TS RemoteApp Manager.
4. **Click** the Add TS RemoteApp Programs link.
5. When the RemoteApp Wizard launches, **click Next**.
6. At this point you will see a list of all of the applications that are available through the Terminal Services. Select the check box that corresponds to the application that you want to make available through TS RemoteApp.
7. **Click Next**.
8. Review the summary information.
9. **Click Finish**.

Making Applications Available to End Users

Once an application has been configured for use with TS RemoteApp, you will need to make the application available to your users. This is accomplished by creating either an .RDP or an .MSI file that you can distribute to your user's desktops.

You would typically use an MSI file if you wanted to roll the remote application out to the desktops by using Group Policy settings or by using a desktop management application such as System Center Configuration Manager. An RDP file is more appropriate for manual deployments or for E-mailing an icon to users that they can use to access the remote application.

You can make a remote application available to users by following these steps:

1. **Open** the Server Manager on your Terminal Server.
2. Navigate through the console tree to Server Manager | Roles | Terminal Services | TS RemoteApp Manager.
3. Locate the **RemoteApp Programs** section in the lower portion of the center console pane.
4. **Select** the application that you want to make available through TS RemoteApp.
5. **Click** the Create RDP File link.
6. When the RemoteApp Wizard starts, **click Next**.
7. Specify the location to which you want to save the RDP file that you are creating.
8. Verify that the Terminal Server name, authentication options, and port number are correct.
9. **Click Next**.
10. **Click Finish**.

Terminal Service Web Access

Another component that can be used for providing users with access to remote desktop sessions or remote applications is Terminal Service Web Access, which is officially known as TS Web. TS Web is designed to allow users to access resources that are hosted on a Terminal Server by using their web browser. It is worth noting however, that Remote Desktop sessions and remote applications do not run within the Web browser. Instead, the browser is used as a means for establishing connectivity with the Terminal Server. Once the user clicks on a hosted resource, Windows will open the Remote Desktop Client and use it to launch the remote desktop. Remote applications work in much the same way they would if the TS RemoteApp icon were installed directly on the user's desktop.

Configuring TS Web

As the name implies, the TS Web role makes use of Internet Information Server (IIS). Because of this, the TS Web role is typically only run directly on the Terminal Server in the smallest of environments. In larger organizations, the TS Web role is usually run on either a dedicated server or on an infrastructure server.

Once the TS Web role has been deployed on a server, the TS Web server must be added to the TS Web Access Computers group on the Terminal Server. After doing so, you will have to configure the TS Web server's data source. In other words, you must tell the TS Web server where it will get a list of remote applications.

Finally, the TS Web server is primarily suitable for use in smaller environments. The TS Web Server can be used in enterprise class environments; but, it can be difficult to make this role work with a full-fledged Terminal Server farm.

Making a RemoteApp Available Through TS Web

Remote applications can be made available through TS Web so that users can launch the remote applications through a Web interface. To make a remote application accessible through TS Web, follow these steps:

1. **Open** the Server Manager on your Terminal Server.
2. Navigate through the console tree to Server Manager | Roles | Terminal Services | TS RemoteApp Manager.
3. Locate the remote application of interest at the bottom of the console's center pane.
4. **Right click** on the remote application and choose the Show in TS Web Access command from the resulting shortcut menu.

If you later decide that you do not want the remote application to be available through TS Web, you can right click on the listing for the remote application and choose the Hide in TS Web Access command from the shortcut menu.

Configuring TS Web Access

In some cases, it may be necessary to perform a small amount of configuration on the TS Web Access site before it will function properly. This configuration consists of specifying the TS Web server's data source. In other words, you need to point the TS Web server to your terminal server. You can accomplish this by following these steps:

1. **Click** on the Windows Start button.
2. Choose the Terminal Services | TS Web Access Administration commands from the Start menu.

3. Windows will launch Internet Explorer and bring up an administrative interface for TS Web. When this screen appears, enter the name of your Terminal Server into the Terminal Server Name field. If TS Web Access is running on the same machine as the Terminal Services then the Terminal Server Name should be set to localhost.
4. **Click** Apply.

Application Deployment

In a non-Terminal Server environment, Microsoft's preferred method for deploying applications to the desktop using native Windows Server 2008 capabilities involves using group policy settings. The group policy gives you two different options for application deployment. You can publish applications or you can assign applications.

The difference between publishing and assigning an application has to do with the way in which the application is actually deployed. Most applications are assigned rather than published.

The group policy contains both user specific settings and computer specific settings. As such, an application can be assigned to either a user or to a computer. If an administrator assigns an application to a computer, the application is installed on that computer automatically.

If an application is assigned to a user, the application is not automatically installed. Instead, when the user logs in, Windows will create an icon for the application. The first time the user tries to run the application, the application will be installed; but, the application is not installed until that point.

Publishing an application works differently. When you assign an application, you can assign it to a user or a computer. When you publish an application, you can only publish it to users. As you will recall, assigning an application to a user causes an icon to be created on the user's desktop. When the application is automatically installed the first time the user tries to run the application.

In contrast, publishing an application does not install the application. It merely makes the application available to the user. It is then up to the user to actually deploy the application.

The biggest drawback to using group policy for application deployment is that the Active Directory does not perform any sort of checks to make sure that the target computers adhere to the application's system requirements.

Application Virtualization

Another problem with deploying and managing applications on desktop computers is that some applications are incompatible with each other. For example, you cannot install Microsoft Office 2007 alongside Microsoft Office 2003. While there is rarely a need to run two different versions of Microsoft Office on a single desktop, this scenario illustrates the point that some applications simply will not coexist.

One way to get around application compatibility issues is through the use of application virtualization. Application virtualization is different from server virtualization in that it does not involve creating a virtual machine with a full blown Windows operating system. Instead, application virtualization involves packaging applications in a way that keeps the various application components (such as registry entries and DLL files) isolated from the rest of the operating system. By doing so, it becomes possible to run otherwise incompatible applications alongside one another.

Microsoft's application virtualization solution is a part of the Microsoft Desktop Optimization Pack (which is a part of the Microsoft Software Assurance program).

The interesting thing about Microsoft's application virtualization solution is that virtualized applications are never actually installed to the desktops. Instead, virtualized applications are streamed to the users on an as needed basis.

Application virtualization is also sometimes used in a Terminal Server environment. When an application incompatibility exists the solution in the past was often to silo the application. This means setting up a dedicated Terminal Server whose job it was to host the incompatible application as a RemoteApp. The dedicated server was necessary because the application could not be installed on the same server as contained conflicting applications.

Application virtualization eliminates the need for Terminal Server siloing. Application Virtualization allows multiple applications to be installed on a single Terminal Server regardless of whether or not those applications would normally have compatibility issues with one another.

System Center Configuration Manager

Even though the 70-646 exam focuses on Windows Server 2008, Microsoft does expect you to know what System Center Configuration Manager is and why it is useful.

System Center Configuration Manager 2007 is Microsoft's replacement for the old System Management Server (SMS Server). It is designed primarily for desktop management in enterprise environments. From an application deployment prospective, Microsoft expects you to know that System Center Configuration Manager allows you to set up deployment tasks that can be used to automatically roll out applications to the desktops on your network. It is even possible to set up rules so that only systems that meet the application's minimum system requirements will receive the application that is being deployed.

Another advantage to using System Center Configuration Manager is that it offers reporting capabilities. This means that after you roll out an application to a group of desktops you can generate a report that shows you exactly which desktops received the application and whether or not any problems have occurred. These types of reporting capabilities simply do not exist for Active Directory based application deployments.

Finally, System Center Configuration Manager supports remote management for desktops. This means that if a user contacts the helpdesk for assistance then it is possible for a technician to connect to the user's PC and remotely correct the user's problem. It is possible to remotely administer a user's PC both on the local area network and over the Internet.

Change Management

As networks become larger in size it becomes impossible for a single person to manage the entire network. Most large organizations employ a number of network administrators. All share the administrative duties. This is referred to as decentralized administration.

The key to making decentralized administration work is to practice change management (which is sometimes also referred to as change control). The basic premise behind change management is that any time an administrator makes a change to a server, that change should be documented. This applies to all administrative actions. For example, a patch installation, registry modifications, and hardware replacement are all types of changes that need to be documented.

The reason why change management is so important is because in a hectic IT environment there is simply no way to keep all of the other administrators informed of the administrative changes that you have made. That being the case, you should document all changes that have been made, as well as when the change took place, why the change was made, and by who.

Having solid change management documentation in place can greatly help with the troubleshooting process when problems occur. If a server suddenly begins to experience problems, it is easy to look back at the change management document for that machine to see if any changes have recently been made that might have triggered the problem.

Application Maintenance

One of the big challenges that network administrators face is keeping applications up to date. Today almost every application vendor publishes updates to their applications. Some vendors might release an application update once every few months, while others release multiple updates each week.

One of the big problems with keeping applications up to date is that there is no universal standard for deploying application updates. For the most part, each application vendor has their own method of delivering updates for their application. This can make it especially tough for network administrators to keep track of which updates have been installed and where.

A lot of applications take a 'set it and forget it' approach to application updates. In other words, the application can be configured to download and apply updates automatically. At first, it might seem as if this approach to update delivery might make an administrator's life easier, but there are a couple of problems with using this approach.

One potential problem is that if every application is configured to download and apply updates automatically, the update process can have a tremendous impact on the available Internet bandwidth. Imagine, for example, that you have a thousand desktops, each of which independently download the same update. If the update was very large or if the update process occurred on a frequent basis, the Internet connection could quickly become congested.

The other problem with allowing applications to download updates automatically is that occasionally application vendors publish buggy updates. As such, an administrator needs to be able to test an update before rolling it out to the entire organization in an effort to ensure that the update is not going to break anything.

The Windows Server Update Service

Microsoft offers a couple of different solutions for dealing with these problems. One solution is to deploy updates by using the Windows Server Update Service (WSUS). WSUS is a free server role that you can use to manage the update process for Microsoft products. The server role downloads updates as they become available and then gives the administrator the ability to deploy the updates to the end users once the updates have been tested. The nice thing about using this approach is that it helps to conserve Internet bandwidth because only one copy of each update is downloaded, rather than a separate copy for each PC.

WSUS can be configured to operate in a centralized or in a decentralized manner. When WSUS is configured for centralized patch management, all of the updates are downloaded to one centralized WSUS server. Once an administrator approves those updates they are then pushed out to the other computers on the network.

A decentralized WSUS deployment is known as a distributed configuration. Distributed configurations are used when patch management needs to be handled separately for different parts of the network. In these types of situations, a central WSUS server downloads the updates, but each of the areas that need to be managed also has its own WSUS server. For example, an organization that has a lot of branch offices might place a separate WSUS server in each branch office. These servers pull the updates from the central WSUS server; but, the approval and distribution process is handled by the individual WSUS servers.

In larger organizations it is possible to use multiple WSUS servers as a way of providing redundancy. WSUS is compatible with Windows Server 2008 services such as Network Load Balancing (NLB), the Distributed File System (DFS), and failover clusters.

Network load balancing can be used in larger organizations to prevent a single WSUS server from having to supply updates to all network clients. Without NLB, the WSUS server could potentially become overwhelmed.

The reason why DFS should be used with WSUS has to do with the sheer volume of updates that Microsoft releases. As you know, the central WSUS server is responsible for downloading all of the latest updates. Although it is possible to store the downloaded updates locally, doing so can be a bad idea.

It's easy to think of the WSUS server as only containing the most recent updates that need to be deployed, but the WSUS server may actually store many gigabytes of update files. Even if all of the computers on your network are up to date, older updates may still be needed. If you were to bring a new machine online for example, Windows will have to determine what updates (if any) have already been applied to the new machine and must then pull all of the remaining updates from WSUS. This process obviously involves far more than just the latest updates.

So what does all of this have to do with DFS? Imagine if your WSUS server were to fail. If the updates were stored locally on the failed DFS server then all of the updates would have to be re-downloaded. This means tying up your Internet connection while a huge amount of data is downloaded and you won't be able to apply any more patches until everything is fixed. Using DFS can help administrators to avoid having to re-download patches in the event of a server failure because patches can be replicated to multiple DFS servers throughout the enterprise.

WSUS can also take advantage of the Failover Cluster service. Even though WSUS itself is not an application that can be clustered, WSUS does depend on Microsoft SQL Server. It uses a SQL Server database to keep track of which updates have been applied to which machines. Even though WSUS itself can't be clustered, SQL Server can be.

Although WSUS solves many of the problems associated with application management, it does have one major drawback. WSUS can only be used to manage updates for Microsoft products. If you need the ability to centrally manage updates for non-Microsoft products then your best bet will be to use System Center Configuration Manager.

System Center Configuration Manager has the ability to package application updates and then deploy them to the PCs on your network. Of course this only works for updates from vendors who allow you to manually download updates rather than requiring updates to be downloaded directly through the application.

In some cases it is also possible to deploy software updates through Group Policy. The catch is that you can only use Group Policy to deploy .MSI files. If software updates are published in another format (such as an EXE file or a ZIP file), you will have to repackage the application as an MSI file before you will be able to assign the update to your desktop computers using group policy settings.

Perhaps the biggest potential pitfall to deploying application updates using Group Policy settings is that the group policy does not have a way of checking to see what is currently installed on each PC. Therefore, it is possible to accidentally assign an update to a PC that doesn't even have the application that is being updated installed. Never mind that group policy lacks the ability to report on the delivery of the update.

Maintaining Windows Server 2008

As is the case with any server operating system, there are certain management tasks that must be performed on servers that are running Windows Server 2008. This can include tasks such as defragmenting storage volumes or applying software updates.

Although server maintenance can be performed manually, doing so is usually impractical due to the number of servers in the organization. Even small organizations may have dozens of Windows 2008 servers because of the proliferation of virtual machines. It is therefore important to understand how to manage Windows Server 2008 in large scale deployments.

Management Tools

Microsoft has introduced several new management tools in Windows Server 2008 that did not exist in previous versions of Windows Server. One such tool is the Server Manager. As you have already seen throughout this book, the Server Manager allows for the installation, configuration, and management of server roles.

ServerManagerCMD

Microsoft has also included a command line version of Server Manager. The command line version, which is known as ServerManagerCMD, allows you to manage roles and features from the command line.

There are X main switches that can be used with this utility. The primary switches include:

- **Query** – Displays a list of all of the roles, role services, and features that are available for installation on the server.
- **InputPath** – Allows you to specify an answer file to be used during a role installation.
- **Install** – Used to install a role, role service, or feature.
- **Remove** – Used to remove a role, role service, or feature.
- **Help** – Provides the full syntax of the command. The syntax is also available from Microsoft
- **Version** – Displays the Server Manager version number.

There are a number of options that can be used with the various switches, but the basic syntax is very simple. For example, if you wanted to see which roles, role services, and features were currently installed and which were available for installation, you would use this command:

```
ServerManagerCMD -Query
```

Likewise, if you wanted to install the Web Server role you could accomplish the task by using this command:

```
ServerManagerCMD -Install Web-Server
```

Windows PowerShell

Another new management tool that is available in Windows Server 2008 is Windows PowerShell. Windows PowerShell is a scripting environment that is completely separate from the Command Prompt window that has always existed in previous versions of the Windows operating system.

Windows PowerShell is extremely powerful and can be used to script virtually any management operation. Furthermore, almost all of the newer Microsoft server products are designed with PowerShell integration. In fact, products such as Exchange Server 2010 are designed so that only the most basic management tasks can be performed through the GUI. All of the more complex configuration changes must be done at the command line through PowerShell. One of the most important new features in Windows PowerShell 2.0, which is included with Windows Server 2008 R2, is remote execution. You can execute PowerShell cmdlets against remote machines and see the results on your local console.

You can install Windows PowerShell by following these steps:

1. **Open** the Server Manager.
2. **Select** the Features container.
3. **Click** on the Add Features link.
4. **Select** the Windows PowerShell option.
5. **Click** Next.
6. **Click** Install.
7. When the installation process completes, **click Close**.

Remote Management

Often, when administrators perform server management tasks, they do so remotely rather than sitting down at the server console. Although the Server Manager does a good job for managing servers locally, it lacks the ability to effectively manage remote servers. Many of the Server Manager's tools and role management utilities simply cannot be connected to a remote server.

When you install a role through Server Manager, Windows usually creates a dedicated tool for administering that role and places it on the server's Administrative Tools menu. For the most part, the administrative tools are identical to the management tools found within Server Manager; but, almost all of the administrative tools support local and remote administration.

Normally the administrative tools are only installed once the role that the tool is designed to manage is installed. For example, the DNS management console would not appear on the Administrative Tools menu unless the server had been configured to act as a DNS server.

For those who wish to remotely manage servers, Microsoft offers the Remote Server Administration Tools (RSAT). The Remote Server Administration Tools are a collection of administrative tools that can be installed onto a Windows 2008 server regardless of the roles that the server is running. This makes it possible to manage remote servers without having to install unnecessary roles on your management server. Incidentally, RSAT can also be downloaded and installed on Windows Vista and Windows 7.

The RSAT tools are especially helpful for managing servers that are running Server Core configurations. As you may recall, Server Core deployments do not have a full-fledged GUI interface. That being the case, your options for managing a Server Core machine include using the command prompt or remotely managing the machine using the RSAT tools or something similar.

Installing the Remote Server Administration Tools

The Remote Server Administration Tools are included with Windows Server 2008 and are available for download for Windows Vista and Windows 7. You can install the Remote Server Administration Tools on Windows Server 2008 by following these steps:

1. **Click** the Start button.
2. **Open** a Command Prompt window.
3. **Type** `ServerManagerCMD -install rsat -<the tools that you want to install>`.
4. When you are done, the tools that you have installed will appear on the server's Administrative Tools menu.

Remote Desktop

Another way to manage remote servers is through Remote Desktop. Remote Desktop allows you to access a remote server's desktop through an RDP session. The underlying architecture works similarly to the Terminal Services except that you are not required to install the Terminal Services role.

In an effort to prevent organizations from circumventing the Terminal Services licensing requirements, Microsoft limits each Windows Server 2008 machine to using two concurrent Remote Desktop sessions.

Enabling Remote Desktop

Remote Desktop is not enabled by default. You can enable Remote Desktop on a Windows Server 2008 machine by following these steps:

1. **Click** the Start button.
2. **Right click** on the Computer container.
3. **Choose** the Properties command from the shortcut menu.
4. When the system properties sheet appears, **click the Remote Settings link**.
5. **Choose** to allow connections from remote computers. You can either require that the remote machines use Network Level Authentication, or you can disable Network Level Authentication. Requiring Network Level Authentication is good for security, but limits your ability to connect from older network clients.
6. **Click** Select Users.
7. **Verify** that the Administrators group was added automatically.
8. Add any additional users or groups that might be required.
9. **Click** OK.
10. **Click** OK.

Server Core

The section on deploying Windows Server 2008 discussed how it was possible to perform a server core installation in which Windows Server 2008 is installed without a full-blown GUI. Machines running server core operating systems are not suitable for all tasks; but, they can be used for infrastructure services such as DNS, DHCP, file services, and print services.

Because server core operating systems only offer a command prompt interface, they are usually managed remotely. Even so, there are two commands that you will need to be familiar with for the 70-646 exam. These commands include: **OCLIST.EXE** and **OCSETUP.EXE**.

The OCLIST command shows you which roles are installed on Windows. Simply enter the OCLIST command within a Command Prompt window and Windows will display a list of all of the server roles as well as which of those roles are currently installed on the server.

The OCSETUP command is used to install and uninstall server roles on a server core machine. The full command syntax is as follows:

```
OCSETUP </? | /h | /help>
OCSETUP <component> [/uninstall] [/passive]
[unattendfile:<file>] [/quiet] [/norestart] [/log: <file>]
[/X:<parameters>]

/?, /H, /help - display help information
<component> - name of the component to be installed or removed
[/uninstall] - uninstall the specified component
[/passive] - unattended mode, shows the progress only
[/unattendfile:<file>] - allows you to provide an answer file
[/quiet] - quiet mode with no user interaction
[/norestart] - do not reboot the server
[/log:<file>] - use a non default logging location
[/X:<parameters>] - Allows extra parameters to be passed to
the installer
```

The OCSETUP command is usually prefaced with the START /W command. This command causes the OCSETUP command to execute within a new window. The advantage of this is that it frees the existing command prompt window for other tasks while the installation is being performed. More importantly however, the /W switch keeps the window from being closed automatically by OCSETUP.

To give you a better idea of how all of this works, you can use the following command to deploy the DHCP Server role:

```
Start /W OCSETUP DHCPserverCore
```

Server Security

Most Microsoft certification exams place a heavy emphasis on security; Exam 70 – 646 is no exception. Windows Server 2008 contains several different security features that you need to be aware of prior to taking the exam.

The Windows Firewall

One security feature that you can expect to be tested on is the Windows Firewall. Windows Server 2008 uses the same firewall as is found in Windows Vista. This firewall, which is enabled by default, allows for granular control over inbound and outbound connections to and from the server. Application servers use firewall rules to open any ports that are necessary for the applications to be able to function.

The Windows Firewall can be configured through the server's Control Panel; but, it is more common to configure the firewall settings through group policies.

Auditing

It is easy to think of auditing as a mechanism for performing forensic investigations after a security breach has occurred. However, auditing is also useful for change management, especially when it comes to the Active Directory.

In Windows Server 2003, changes to the Active Directory could be audited, but only the new values were recorded within the audit logs. This was somewhat of a problem because if an administrator ended up accidentally setting an Active Directory attribute to use an incorrect value then there was no way to know what the previous value was unless the administrator happened to either remember the previous value or to have written it down. Windows Server 2008 differs in its Active Directory auditing in that when a change is made to an Active Directory object, both the old and the new values are recorded within the audit logs.

Microsoft has made auditing enhancements in Windows Server 2008 R2. One such enhancement is Global Object Access Auditing. This feature allows administrators to define system access control lists for either the registry of the file system. These system access control lists can be applied to every object of a specific type.

Another new auditing feature is the Reason for Access Reporting feature. When a user attempts to access an object, the Reason for Access reporting feature can document the permission that either gave the user access to the object or denied the user access to the object.

Finally, previous versions of Windows used nine basic types of auditing. These nine auditing types have been replaced by fifty three granular auditing types. This allows administrators to audit exactly what they need to without performing excessive auditing.

Delegation of Administration

As your network grows it may eventually become impossible to effectively manage the entire network by yourself. As stated earlier, larger organizations employ a number of different network administrators all of whom are collectively responsible for managing the network through decentralized administration. In some organizations each administrator has full control over the entire network. More often however, some administrators specialize in specific areas. For example, an organization might employ an administrator to be responsible for managing the Active Directory. This type of administrator would need full rights to the Active Directory, but probably would not require full rights to things like DFS targets or WSUS.

Because there may be a need to provide certain employees with limited administrative capabilities, Windows Server 2008 makes it possible to delegate administrative tasks. This means that it is possible to assign network administrators the rights that they need in order to do their jobs, but without giving them excessive rights.

Any time that you grant any sort of permissions to anyone, you should always apply those permissions to a group, not to individual users. This holds true regardless of whether you are delegating administrative authority or just casually giving users access to a file share. Assigning permissions to users rather than to groups can be incredibly difficult to track down everything that a specific user has rights to.

Another reason why permissions should be delegated on a group basis has to do with object ownership. Imagine, for instance, that you were to give a user the ability to create Active Directory objects. If you later revoked that ability, the user would still have full control over every object that they had created because they are the owner of the object. On the other hand, if you delegate the authority to a group, when a group member creates an Active Directory object, it is the group not the individual who owns the object. This means that if you remove the user from the group, the user will no longer have administrative control over the Active Directory objects they created.

Active Directory Delegation

One of the most common areas in which administrative control might be delegated is within the Active Directory. For example, an administrator might give the helpdesk staff the ability to reset user's passwords. Likewise, some organizations like to delegate the ability to create user accounts to the HR Department. There are a number of different reasons why a user might be delegated some level of administrative control within the Active Directory.

Active Directory delegation is performed through the Delegation of Control Wizard. The Delegation of Control Wizard provides an interface to delegate the most common Active Directory management tasks.

It is worth noting that although the Delegation of Control Wizard can be used to delegate Active Directory permissions, the Wizard cannot be used to revoke those permissions later on. If you decide that your delegations have gotten out of hand and you need to put everything back to the way that it was, you must go through the Organizational Unit's (OU) Security tab. This tab provides a mechanism for resetting the default permissions within the Organizational Unit.

Performing an Active Directory Delegation

Active Directory delegations are established in roughly the same way regardless of the permissions that are actually being delegated. The procedure below illustrates a method for delegating the ability to reset passwords. The same basic procedure can also be used for performing other types of Active Directory delegations.

To create a new group and delegate permissions for group members to be able to reset Active Directory passwords, follow these steps:

1. **Open** the Active Directory Users and Computers console.
2. **Select** the Users container.
3. **Right click** on an empty area within the list of users and groups.
4. **Select** the New | Group commands from the resulting shortcut menus.
5. When the New Object dialog box appears, enter the name for the group that you are creating.
6. Set the Group Scope to Global.
7. Set the Group Type to Security.
8. **Click** OK.
9. **Right click** on the group that was just created and choose the Properties command from the shortcut menu.
10. When the group's properties sheet appears, select the **Members tab**.
11. **Click** the Add button and add members to the group.
12. **Click** OK.
13. **Right click** on the Users container and choose the Delegate Control command from the shortcut menu.
14. When Windows launches the Delegation of Control Wizard, **click Next**.
15. **Click** the Add button.
16. Enter the name of the group that you just created.
17. **Click** OK.
18. **Click** Next.
19. The wizard will now display a list of common delegation tasks. **Select** the Reset User Passwords and Force Password Change at Next Logon check box.
20. **Click** next.
21. **Click** Finish.

Removing a Delegation

If you want to remove the delegation that you have just created you can do so by following these steps:

1. **Open** the Active Directory Users and Computers console.
2. **Select** the Advanced Features command from the console's View menu.
3. **Right click** on the Users container and select the Properties command from the shortcut menu.
4. **Select** the properties sheet's Security tab.
5. **Select** the group for whom you want to remove the delegation.
6. **Click** Remove.
7. **Click** OK.

Server Management Delegation

The other type of delegation that can be performed within Windows Server 2008 is server management delegation. Server management delegation gives a group of users the ability to manage an individual server. This is completely different from any type of delegation which might be performed at the Active Directory level.

Server management delegation is performed by adding global groups to a server's local groups. Typically, if you need to give a group the ability to manage a server, you will need to add the group to the server's local Administrators group. While making the group members administrators might seem like overkill, this is the minimum level permission that is needed to add roles and features to the server as well as to perform many other server management tasks.

One important aspect of server management delegation that you need to be aware of is that Windows Server 2008 treats domain controllers and member servers differently. When you had a global group to the local Administrators group on a member server, you give the group members the ability to manage the server. In contrast, if you add a global group to the Administrators group on a domain controller then you have effectively granted administrative control over all of the domain controllers in the entire domain. Domain controllers share groups.

This concept further underscores the point that domain controller should be running on dedicated machines (physical or virtual). If a domain controller is also acting as an application server or a file server, then it is impossible to delegate control over the server without also delegating the group control over the other domain controllers within the domain.

Delegating Server Management

You can delegate the management of a member server to a group's members by following these steps:

1. **Open** the Computer Management Console.
2. When the console opens, **right click** on the Computer Management (local) container and select the Connect to Another Computer command from the shortcut menu.
3. When prompted, enter the name of the server for which you want to delegate control.
4. **Click** OK.
5. Navigate through the console tree to Computer Management | System Tools | Local Users and Groups | Groups.
6. **Double click** on the Administrators group.
7. When the Administrator Properties dialog box opens, click the **Add button**.
8. Enter the name of the global group to whom you want to delegate control over the server.
9. **Click** OK.

Group Policies

Group policies are the primary security mechanisms for Active Directory environments. Group policies can be used to apply individual security settings on either a per user or a per computer basis. The per user and per computer policies are applied through links to the site, domain and organizational units within Active Directory.

Group policies have existed in previous versions of Windows Server; but, Microsoft has made some significant changes in Windows Server 2008. One of the most noteworthy changes is that Microsoft has changed the format used by administrative template files.

In Windows Server 2003, administrative templates were composed of .ADM files. Although .ADM files worked relatively well, they had one big problem. Administrative Templates did not work well in organizations in which multiple languages were being used. Administrative templates only supported a single language regardless of any localizations that had been put into place on the network.

In Windows Server 2008, Microsoft has replaced .ADM files with XML based .ADMX files. In addition, each .ADMX file offers support for multiple localizations through the use of language files. The language files use an .ADML extension.

Another thing that you need to know about this new format for group policy templates is that they are only compatible with Windows Server 2008, Windows Server 2008 R2, Windows Vista, and Windows 7. Older versions of Windows are only capable of working with the older style of administrative templates. However, older Windows servers can store Group Policy settings for newer clients, such as Windows 7, in their Active Directory SysVol folders for processing purposes.

The Central Store

Any group policy objects that you create need to be stored within a central store on your domain controllers. The central store is a central repository that gets automatically replicated to all of the domain controllers within the domain. That way, you can make sure that every domain controller has an identical copy of your group policy objects.

The process of creating a central store is really simple. To do so, follow these steps:

1. **Open** Windows Explorer on one of your domain controllers.
2. Navigate through the folder structure to %systemroot%\SYSVOL\domain\Policies.
3. Create a sub folder called **PolicyDefinitions**.
4. **Copy** all of your ADMX files into the PolicyDefinitions folder.
5. **Create** a language folder beneath the PolicyDefinitions folder. The name that you will use for this folder varies depending on the language that you are using. For American English the folder should be called EN-US.
6. **Copy** your ADML files into the language folder.

The Central Store can also be accessed from outside of the domain controller. In doing so, the path to the central store is: \\<your domain controller's fully qualified domain name>\SYSVOL\<your domain name>\Policies\PolicyDefinitions.

Group Policy Hierarchies

The most important thing that you need to understand about group policies is that group policies are applied in a hierarchical and cumulative manner. There are several places throughout the Active Directory and the local computer in which group policy objects can be applied. All of these group policy objects are combined in a way that creates the effect of policy.

Group policy objects can be applied on a per user and/or a per computer basis. It is very common to have both user level and computer level policies so that you can control what happens if certain users logon to certain machines.

The group policy hierarchy is processed in the following order:

1. The local computer group policy object is applied.
2. Active Directory site level group policy objects are applied.
3. Domain level group policy objects are applied.
4. Organizational unit group policies are applied.

As group policy objects are applied, the settings within those group policy objects are cumulative. In other words, a setting that is applied at the local computer level will remain in effect regardless of what other policies are applied unless a higher level policy (such as one applied at the site or domain level) overwrites that policy setting.

When it comes to policy settings, the settings that is applied most recently takes precedence. Therefore, if two group policy objects contain contradictory settings, the setting within the higher level policy will be the one to take effect. If a lower level policy contains a setting that is not addressed in any of the higher level group policy objects then that setting will remain in effect and become a part of the effective policy.

Resultant Policies

Because the effective policy can contain user and computer elements found in group policy objects scattered throughout the Active Directory and the local computer, it can be a little bit complicated to figure out what policy settings are going to be in effect within a given situation. Occasionally, a policy setting may end up being in effect unexpectedly. In the past, unexpected policy settings had to be tracked down manually. In Windows Server 2008 however, Microsoft gives you tools for determining where and effective policy setting came from. You can track down a group policy setting by following these steps:

1. Log on to a domain controller.
2. **Select** the Group Policy Management command from the server's Administrative Tools menu.
3. Navigate through the console tree to Group Policy Management | <Your Forest> | Group Policy Results.
4. **Right click** on the Group Policy Results container and select the Group Policy Results Wizard from the shortcut menu.
5. When the Group Policy Results Wizard begins, **click Next** to clear the Welcome screen.
6. Specify the computer for which you want to test the group policy settings (remember that group policies are applied on a per computer and per user basis).
7. **Click Next**.
8. Specify the user for whom you want to test the group policy settings.
9. **Click Next**.
10. **Click Next** to clear the summary screen.
11. **Click Finish**.

The results are presented in a container beneath the Group Policy Results container. This container will have a name that reflects the test. For example it might be called Administrator on DC1.

When you select this container, the results pane will display three tabs. The Summary tab will allow you to see what policy settings are in effect. The Settings tab will show you the individual policy settings that are in effect and which Group Policy Objects the various settings were found in. The Policy Events tab shows Event Viewer logs related to the group policy settings.

Backing Up Group Policy Objects

Once you have created any necessary group policy objects, it is a good idea to back them up. That way, if you ever make any changes to your group policy objects and later discover that those changes are causing you problems, you can revert to a previous version of the group policy object by simply restoring a backup.

Backing up a group policy object can also be useful during the initial planning and creation phase. Many organizations use many different group policy objects that are placed in various levels of the Active Directory hierarchy. If you have certain settings that should be consistent across all of your group policy objects, you can create a group policy object that has those settings in place and then back it up. You can then use that backup as a means for creating any additional group policy objects that you may need.

It is worth noting that if you ever have to restore a group policy object, the restored group policy object completely overwrites the group policy object that it is replacing. Windows will not merge the group policy settings from the version that you are restoring with any settings within the existing policy. Restoration of group policy objects is an all or nothing process.

To back up your group policy objects, follow these steps:

1. **Open** the Group Policy Management Console.
2. Navigate through the console tree to Group Policy Management | <your forest> | Domains | <your domain> | Group Policy Objects.
3. **Right click** on the Group Policy Objects container.
4. **Choose** the Back Up All command from the shortcut menu.
5. When the Back Up Group Policy Object wizard begins, specify a location to store the backup.
6. Enter a description for the backup that you are creating.
7. **Click** the Back Up button.
8. **Click** OK.

Restoring a Group Policy Object

If you need to restore a group policy object you can do so by following these steps:

1. **Open** the Group Policy Management Console.
2. Navigate through the console tree to Group Policy Management | <your forest> | Domains | <your domain> | Group Policy Objects.
3. **Right click** on the Group Policy Objects container.
4. **Select** the Manage Backups command from the shortcut menu.
5. When the Manage Backup window appears, select the Group Policy Object that you want to restore.
6. **Click** the Restore button.
7. **Click** OK.
8. **Click** Close.

Monitoring Servers

In any organization it is important to be able to monitor your servers to ensure that they are healthy and that there are no issues which may be affecting the end users. For server monitoring in an enterprise environment, Microsoft's monitoring tool of choice is Systems Center Operations Manager. For those whose budgets do not support purchasing Systems Center Operations Manager licenses however, Microsoft does include several different monitoring tools within Windows Server 2008.

The Event Viewer

Although the Event Viewer has been a part of Windows Server since the very beginning, it has been completely overhauled the Event Viewer in Windows Server 2008. One of the nice new features is that when you open the Event Viewer, it displays an Overview and Summary page. This page shows you a summary of the latest administrative events as well as a summary of the logs. The main purpose of this screen is to make it easy to see any issues that need immediate attention without the administrator having to go looking through the logs.

Another change that Microsoft has made to the Event Viewer is the inclusion of several new types of logs. The Event Viewer still offers the Application, Security, and System logs that were found in previous versions of Windows Server, but they have added a number of new log types that focus on events related to specific areas of the operating system. For example, there is now a Directory Service log and a DFS Replication Log. Additionally, the Event Viewer is extensible so that third party application vendors now have the ability to create application specific logs that are accessible through the Event Viewer.

Probably the most significant change that Microsoft has made to the Event Viewer is the ability to set up event log subscriptions. Event log subscriptions allow for the central monitoring of event log entries across multiple servers. Administrators are able to define the types of events that they want to include in central subscriptions so that only the most critical events are included in the subscription. It is also possible to configure the Event Viewer to run a script in the event that certain log entries are encountered. Such a script could be used to take corrective action or to generate a notification such as an E-mail message.

To create an Event Viewer subscription, follow these steps:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Diagnostics | Event Viewer | Applications and Services | Subscriptions.
3. When you select the Subscriptions container you will see a message telling you that in order to work with subscriptions the server must be running the Windows Event Collector Service. The message goes on to ask you if you want to automatically start the service when Windows is restarted. Click Yes.
4. **Click** the Create Subscription link.
5. When the Subscription Properties dialog box opens, enter a name for the subscription that you are creating.
6. Enter a description for the subscription.
7. Set the Destination Log option to **Forwarded Events**.
8. Set the Subscription type to Collector Initiated so that the server will pull events from other servers.
9. **Click** the Select Computers button.
10. Enter the names of the computers that you want to pull events from.

11. **Click** OK.
12. **Choose** the Edit option on the Select Events button.
13. Specify the types of events that you want to extract from your target servers.
14. **Click** OK.
15. **Click** OK.

A final new feature in the Event Viewer, is the ability to create a scheduled task that is linked to an event. You can indicate that an email notification should be sent to an administrator or a script should run when a specific event is logged into the log files. This allows for improved automation in server environments.

Monitoring Server Performance

One of the complaints that network administrators receive most often from the end users is that their computers are running slowly. In this type of situation, it is helpful to be able to quantify whether or not the server is actually running slowly so that if a problem is occurring you can begin troubleshooting the problem.

Microsoft offers a couple of different tools that you can use to gauge a server's performance. One tool that is often overlooked is the Task Manager. The Task Manager is excellent for taking a quick look at the server's current performance. When you open the Task Manager simply go to the Performance tab. This tab contains a CPU graph and a memory graph that you can use to assess the machine's current performance.

Although the Task Manager is great for quickly assessing a machine's current performance, it provides a very limited amount of performance data. If you need a little bit more information than what the Task Manager is able to provide, you might find another performance tool helpful.

In Windows Server 2008, Microsoft introduced a new performance monitoring tool called the Resource Monitor. You can access the Resource Monitor directly through the Task Manager by going to the Performance tab and clicking the Resource Monitor button.

The Resource Monitor is designed to provide a deeper level of performance monitoring than what you can get through the Task Manager; but, without providing the technical depth of the Performance Monitor. You can think of the Resource Monitor as a light version of the Performance Monitor. This tool provides performance information on four main areas of the system including CPU, memory, network, and disk. Additionally, the Resource Monitor allows you to analyze the wait chain for a stalled process. If you have ever had a process stall and display the message that reads "Not Responding," you will find this feature useful. It can inform you of the cause of the stall. Simply right-click on the process in the Resource Monitor and select Analyze Wait Chain.

For those situations that require access to more in-depth performance information, the tool of choice is the Performance Monitor. Like the Task Manager and the Resource Monitor, the Performance Monitor can provide a live snapshot of how the system is performing right this minute. However, there are a few things that set the Performance Monitor apart from the Task Manager and the Resource Monitor.

One of the most significant differences that set the Performance Monitor apart is that it can be used to track a system's performance over time. This is important because the only way that you can truly know whether or not a system is experiencing performance problems is to know how the system is supposed to be performing. The best way to accomplish this is to use the Performance Monitor to examine the system's performance when it is running well. You can use this information to establish a performance baseline. You can then compare the system's performance over time or at times when the system seems to be responding slowly to the baseline that was established earlier. By doing so, you can see exactly how the system's performance has changed. Simply knowing which areas of the system are being impacted by the performance problems can go a very long way toward helping you to troubleshoot the problem.

The Performance Monitor can also be extremely useful for capacity planning. For example, you can use the Performance Monitor to monitor a system's disk consumption rate over time. By doing so, you can estimate when the server is likely to be and running low on disk space. This can help you to budget for the additional storage that will be required long before you actually need to make the purchase.

Like most of the other administrative tools that come with Windows Server 2008, the Performance Monitor can be used to assess the performance of either the local system or a remote server. This can be especially helpful if you have servers that are running a server core operating system since those servers do not have performance monitoring capabilities of their own.

The Reliability Monitor

One more performance related tool that you need to be aware of is the Reliability Monitor. The Reliability Monitor is designed to show you how reliable a server has proven to be over time. The reason why this tool is beneficial is because often times network administrators are asked to examine servers with which they may not be immediately familiar. The Reliability Monitor provides an instant picture of the types of issues that the server has experienced over time as well as clues to what might be going on with the server right now.

The Reliability Monitor works by examining performance logs and event logs. It uses this information to calculate a daily reliability score. This score is plotted on a graph so that the administrator can see the system's reliability over time. The reliability score is affected by things like application installations, hardware changes, configuration changes, and system crashes. In essence, the Reliability Monitor is a great tool for determining whether or not changes have recently been made to a server.

The Performance Monitor

As previously mentioned, one of the main benefits to using the Performance Monitor is that it can provide more in-depth information than the other performance tools are capable of providing. When you use the Performance Monitor you have to know exactly what aspects of the system's performance you want to monitor. This is important because the Performance Monitor can track hundreds, if not thousands of different aspects of the server's performance, so you need to know exactly what it is that you're looking for.

In order to measure some aspect of the system's performance, you will have to add a performance counter to the Performance Monitor. A Performance Counter is a mechanism for tracking one very specific aspect of the server's performance. Performance counters are organized into performance objects. You can think of a performance object has a category for system performance. Performance counters that measures similar aspects of the system's performance are grouped into common performance objects. For example, performance counters that are designed to track CPU performance are found beneath the Processor performance object.

Performance counters are extremely granular. Each counter monitors one very specific aspect of system performance. In the real world, it is rare that you would only use a single counter when troubleshooting or monitoring system performance. Most of the time, you would use a combination of several different counters so that you can get a more accurate picture of what is happening with your server.

The Performance Monitor does allow you to manually add counters any time that you want to. This is fine for an impromptu monitoring session; but, if you have certain counters that you monitor on a regular basis, you are usually going to be better off setting up a data collector set.

Simply put, a data collector set is nothing more than a collection of performance counters that you intend to use collectively. Windows Server 2008 provide some data collector sets for you, but it also allows you to set up user defined data collector sets, which is a great option if you have performance counters that you need to monitor on a regular basis.

Data collector sets can be used manually, but you can also configure data collector sets to collect performance data on a scheduled basis. By doing so, you can collect performance information over time for a predetermined set of performance counters. For example, you might configure a data collector set to collect performance data for fifteen minutes each day.

The Windows System Resource Manager

Sometimes being able to monitor a server's performance simply isn't enough. Ideally, administrators would like to be able to control how a server performs. While it is impossible to exercise total control over every aspect of a server's performance, Windows Server 2008 does offer a tool which allows you to ensure that certain types of application servers do not become bogged down during periods of heavy activity.

This tool is called the Windows System Resource Manager. The basic idea behind this tool is that it allows you to control memory and CPU utilization for the different processes on the system. This can be especially helpful in any type of environment in which users are sharing the resources on an application server.

One example of this is that in a Terminal Services environment you might occasionally run into a situation in which one user is using a Terminal Service session in a way that deprives other users of system resources. For instance, suppose that you have 10 users log into a Terminal Server and that nine of those users are using their sessions for word processing. Now let's suppose that the 10th user is running a high demand database application. If left unchecked this user's actions can actually slow things down for the other nine users.

The Windows System Resource Manager helps to alleviate this problem by allowing you to establish certain policies that dictate how system resources are to be used. Rules can be based around processes, users, or sessions. For example, you might create a rule that allocates an equal amount of CPU time to each process that is running on a server.

Of course some users might still consume more system resources than others by running multiple processes. If this is an issue in your environment then you might create a policy that assigns an equal amount of CPU time to each user that is connected to the server. This can keep a single user from monopolizing the server's CPU resources.

In a Terminal Services environment you can also base policies around sessions. For example, you could create a policy that allocates equal amounts of CPU resources to each server session.

Keep in mind that these types of techniques might only be effective if the server workload remains relatively uniform. Sometimes business operations cause servers to have to juggle different workloads at different periods of time. For example, the Finance Department might run a payroll report on the second Friday of every month.

In these types of situations it is possible to use calendar rules to create a schedule for different resource policies. That way, you can make sure that when the payroll report is run it doesn't deplete the server resources. On the other hand, since everyone probably wants to get paid you could configure the resource policies to ensure that the payroll processing received preferential treatment.

Although the Windows System Resource Manager is great for ensuring that system resources are used in a balanced manner, it also has another use. The Windows System Resource Manager can also be used to profile a system's performance.

Suppose for instance that you had a system that routinely ran low on system resources. You could configure the Windows System Resource Manager to monitor resource consumption on the server. By doing so, the Windows System Resource Manager would build a log file that details the exact resources that are being used by each process. That way, you could gain a clear understanding of how the system's resources are actually being used. This technique can be very effective in locating problems such as memory leaks.

The Certificate Services

In Windows Server 2008 the Certificate Services provide the digital certificates that are used by many other security related features. The Certificate Service's official name is the Active Directory Certificate Services. However, you do not have to have an Active Directory environment in order to use the certificate services. There are however benefits to creating a certificate authority within an Active Directory environment.

The Certificate Service's job is to create a Public Key Infrastructure (PKI). PKI is based on the use of certificate authorities whose job it is to issue certificates that can be used for encryption and for confirmation of a computer's identity.

The Certificate Services are made up of a number of different components. The components that are available for use vary depending on the edition of Windows Server 2008 that is being used.

There are two Certificate Services components that can be installed in Windows Server 2008 Standard Edition. The first of these components is the Certification Authority. The Certification Authority is the core component of the Certificate Services that handles the actual certificate distribution process.

The other component that can be installed on servers running Windows Server 2008 Standard Edition is the Web Enrollment component. Web Enrollment refers to a dedicated Web site that users can log into and request certificates from the Certification Authority.

Windows Server 2008 Enterprise Edition offers the same Certificate Services components as the Standard Edition does, but also offers a couple of additional components. The first of these components is known as the Online Responder.

A Certificate Authority maintains a special list called a Certificate Revocation List (CRL). The CRL is essentially a list of any certificates that have expired or that have been revoked. When Windows needs to check the validity of a certificate, it normally contacts the Certificate Authority, downloads the CRL, and then checks the downloaded copy for any references to the certificate that needs to be validated. As you can imagine, this process takes some time. However, when the Online Responder is in use it becomes possible for clients to check a single certificate against the CRL without having to download the entire list.

The other component that is available in Windows Server 2008 Enterprise Edition is the Network Device Enrollment Service. This service makes it possible to issue certificates to hardware devices so that those devices can benefit from the PKI infrastructure that you have in place.

Types of Certificate Authorities

Windows Server 2008 allows for the creation of various types of Certificate Authorities. The two main types of Certificate Authorities that you can create are Stand-alone Certificate Authorities and Enterprise Certificate Authorities.

Standalone Certificate Authorities are certificate authorities that are not connected to the Active Directory.

These types of certificate authorities are useful in situations in which you need to issue certificates to users who exist on a partner network or on an extranet and do not have an Active Directory account on your network.

An enterprise certificate authority is integrated with the Active Directory. There are a number of benefits to using enterprise certificate authorities. One such benefit is auto enrollment. In other words, the computers on your network will automatically request and be provided with a public key certificate. The computers must simply be domain members.

Another choice that you will have to make involves whether you want to install a root certificate authority or a subordinate certificate authority. In an Active Directory environment the first certificate authority that you set up should always be a root certificate authority.

In larger organizations a single certification authority might be inadequate for handling all of the certificate requests. You might also want to distribute certificate authority servers throughout your network in strategic locations (such as branch offices). In these types of situations you would create a certificate authority hierarchy. The root certificate authority sits at the top of the hierarchy and subordinate certificate authorities exist beneath the root.

In high security environments, it has become a common practice to create a root certification authority then take it offline. That way, the root certification authority can never be compromised. In these types of environments, it is the subordinate certificate authorities that do all of the work.

Planning a PKI Hierarchy

Even though it is relatively simple to set up a Certificate Authority, the process is not something to take lightly. There are any number of Windows services which depend on the use of certificates. Therefore, if you implement a poor design there can be major consequences down the road.

The first lesson for implementing a certificate authority is that the root certification authority is critically important. If someone compromises your root certification authority, they own your network. You must, therefore, protect the root certification at all costs. That being the case, you should install the root certification authority onto a dedicated server. Although it might be tempting (from a security prospective) to install the Certificate Services onto a server core deployment, doing so is not supported. The Certificate Services require the .NET Framework, which cannot be used in a Server Core environment.

As previously mentioned, in larger environments it is a good idea to use an offline root certification authority. To do so, you would create a standalone root certification authority and an enterprise subordinate certificate authority. You would then take the root certification authority offline so that it cannot be compromised. You only need to bring the root certification authority online if you decide to create an additional subordinate certificate authority (because the root has to issue a certificate to the subordinate). The rest of the time the root certification authority can remain safely offline.

Installing an Enterprise Certificate Authority

To install an enterprise certificate authority, follow these steps:

1. **Open** the Server Manager.
2. **Select** the Roles container.
3. **Click** on the Add Roles link.
4. When the Add Roles Wizard begins, **click Next**.
5. **Select** the Active Directory Certificate Services check box.
6. **Click** Next.
7. **Click** Next to bypass the summary screen.

8. **Select** the individual services that you want to install.
9. If necessary, click the **Add Required Role Services** button to install IIS.
10. **Click Next**.
11. **Choose** whether you want to create an Enterprise or a Standalone certificate authority.
12. **Click Next**.
13. Decide whether you want to create a Root CA or a Subordinate CA.
14. **Click Next**.
15. Decide if you would like to create a new private key or if you would like to use an existing private key.
16. **Click Next**.
17. **Choose** the encryption algorithm that you want to use. The default is usually acceptable.
18. **Click Next**.
19. Enter a name for your new certification authority.
20. Make note of the distinguished name that Windows automatically generates.
21. **Click Next**.
22. **Choose** a validity period for your certificates. The default validity period is five years.
23. **Click Next**.
24. **Choose** the path where the certificate database should be stored.
25. **Click Next**.
26. If you have chosen to install the Web Enrollment component you will now see a screen introducing you to Internet Information Server (IIS). **Click Next** to bypass this screen.
27. Windows will ask you which IIS components you want to install. **Click Next** to accept the defaults.
28. Take a moment and verify that the information on the summary screen is correct.
29. **Click Install**.
30. When the installation completes, **click Close**.

Network Access Control

Network Access Control refers to a group of technologies that control access to your network based on policies. These policies can regulate access to both the wired and the wireless network. Network Access Control has been implemented in Windows Server 2008 through the Network Policy and Access Services Role.

There are several different components that make up the Network Policy and Access Services role. The component that is probably the most familiar is Routing and Remote Access. This is the service that allows Windows Server 2008 to host a Virtual Private Network (VPN) connection. The Routing and Remote Access service can also be used to facilitate dial in access to the network.

Another component of the Network Policy and Access Services is the Network Policy Server (NPS). The Network Policy Service replaces the Internet Authentication Service (IAS) as Microsoft's RADIUS server. It provides authentication and authorization to users who are trying to connect to the network. The Network Policy Server is the main component of the Network Policy and Access Services.

A third component that is included in the Network Policy and Access Services is the Health Registration Authority. The Health Registration Authority is a certificate based component that allows a server to check to see if computers that are attempting to connect to your network are healthy. As a network administrator you get to decide what constitutes a healthy PC. For example, you might define a healthy PC as a PC that has the Windows firewall enabled and that is running up to date antivirus software.

One last component of the Network Policy and Access Services is the Host Credential Authorization Protocol. The Host Credential Authorization Protocol allows for integration with Cisco's Network Access Control Server device, which performs the same basic functions that allows Windows Active Directory to act as the authenticating source for access to the network.

Network Access Protection

Network Access Protection (NAP) is the main component of the Network Policy and Access Services, and is responsible for controlling who is and is not allowed to access the network. NAP is actually able to examine the client computer to make sure that it adheres to the corporate security policy. For example, you can look at things like whether or not Windows Update is enabled, whether or not antivirus software is installed, and whether or not your antivirus software is up to date.

As powerful as NAP is, it does have one major drawback. Network Access Protection cannot work by itself. It requires the Network Policy Server and the client computer to work together. The reason why this is a problem is because NAP is a new feature for Windows Server 2008 and is therefore not backward compatible with older client operating systems. NAP is only supported by Windows XP SP3, Windows Vista and Windows 7.

In order for client computers to work with NAP, you must enable a NAP Enforcement Agent on the client computers. The NAP Enforcement Agent makes the client computers NAP aware. The NAP Enforcement Agent is typically enabled through group policy settings.

Because NAP is designed to control access to the network, it usually resides at the network entry point. In most cases, this means that a VPN server will be NAP enabled, although it is possible to enable NAP for other types of connections such as for a wireless network. In fact, NAP can control access through VPN and 802.1x. It can also be tied to your DHCP servers so that client computers will not receive an IP address unless they adhere to the Network Access Policy.

In an environment in which a client attempts to connect to the network through a NAP enabled VPN server, the process begins when the client connects to the VPN server and sends a set of authentication credentials in an effort to connect to the network.

When the VPN server receives the user's request, they must then reply back to the client computer stating that network access will be granted, so long as the client computer can prove that it is healthy.

At this point, the client computer assembles what is known as a Statement of System Health (SoH). At a minimum, the Statement of Health includes the client's compliance status, the client's site, and a time stamp reference that is used to identify the NAP policies that the client used to evaluate its compliance.

The Statement of System Health is sent to the VPN server, which then forwards it to the Network Policy Server. The Network Policy Server compares the Statement of System Health against the various network health policies to make sure that the client computer is in compliance with those policies. Upon making this comparison the Network Policy Server categorizes the client computer as compliant, non-compliant, or non-compatible.

As you would expect, a compliant computer completely adheres to the network health policies. A non-compliant computer has some deficiencies which keep it from being compliant with the network health policies. A non-compatible computer is a computer that simply isn't compatible with Network Access Protection; so the Network Policy Server is unable to determine whether the computer is compliant or not.

If a computer is determined to be non-compliant or non-compatible, it is usually placed on a restricted network segment where it can be isolated from the rest of the network. This restricted segment can contain a remediation server. The remediation server's job is to help the computer to become compliant with the network health policies. There are any number of ways in which this can be accomplished. For example, the restricted segment might contain a Windows System Update Server which can apply the latest patches to the computer. Likewise, the remediation server could potentially contain a file share with the latest anti-virus definitions. It is even possible to automatically enable the firewall on the client computer.

After the client has been remediated, it is given access to the corporate network just as if it had been compliant from the very beginning.

Enabling Network Access Protection

The technique that is used to enable NAP varies depending on how you plan to use it. One of the easiest ways to set up NAP is to enable it through DHCP. You can do so by completing the steps below.

This technique assumes that the Network Policy and Access Services role is already installed on the server. You should also have a Windows Server 2008 machine that is configured to act as a DHCP server. This server should already be configured to with a valid scope of IP addresses that it can supply to network clients.

The technique described below causes DHCP clients to be checked to see if they are compliant with the policies that have been set up. If a client is compliant, they will be given an IP address from the existing scope. If the client is found to be non-compliant, they will be connected to an isolated network segment via a router setting. In essence, the DHCP server will use one router IP address for compliant clients and a different router IP address for non-compliant clients. With that said, you can configure your DHCP server for Network Policy Service support by following these steps:

1. **Click** the Start button.
2. **Select** the DHCP command from the server's Administrative Tools menu.
3. When the DHCP console opens, **right click** on the DHCP container and select the Add Server command from the shortcut menu.
4. Choose the DHCP server that you want to connect to and **click OK**.
5. Navigate through the console tree to DHCP | <your DHCP server> | IPv4 | Scope.
6. **Right click** on your scope and select the properties command from the shortcut menu.
7. When Windows opens the scope's properties sheet, select the Network Access Protection tab.
8. **Select** the Enable for This Scope option.
9. **Choose** to use the default Network Access Protection Profile or provide a custom profile.
10. **Click OK**.
11. **Select** the Scope Options container.
12. **Right click** on an empty area in the details pane and choose the Configure Options command from the shortcut menu.
13. When the Scope Options properties sheet opens, select the **Advanced** tab.

14. **Choose** the Default Network Access Protection Class option from the User Class drop down list. This class represents users who are non-compliant.
15. **Select** the Router check box.
16. Enter the IP address of the router that you want to connect non-compliant computers to.
17. **Click** Add.
18. **Click** OK.

When a network client connects to the DHCP server to request an IP address, the DHCP server asks the client for a Statement of System Health. The client computer provides the Statement of System Health to the DHCP Server, which in turn forwards it to the Network Policy Server. The Network Policy Server evaluates the Statement of Health and makes a determination as to whether or not the client is compliant with the network health policies. If the client is compliant then it is given an IP address and directed to a router on the corporate network. Non-compliant computers are sent to a router that's connected to an isolated network segment.

Setting Up Network Access Protection Policies

In the previous section, you saw how to configure a DHCP server so that it only provides IP addresses to computers that comply with the network health policies. However, as of right now, no policies have been created. The technique that you will use to create the various policies varies depending on the network connection method. The technique listed below can be used to create NAP policies for computers that are trying to lease an IP address from a DHCP server:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Network Policy and Access Services | NPS Local.
3. **Choose** the Network Access Protection (NAP) option from the configuration scenario drop down list.
4. **Click** on the Configure NAP link.
5. **Choose** the Dynamic Host Configuration Protocol option from the Network Connection Method drop down list.
6. Provide a name for the policy that you are creating.
7. **Click** Next.
8. The Configure NAP Wizard now asks which RADIUS server will be used to authenticate the user's connection. As you will recall, the Network Policy Server acts as a RADIUS server. Therefore, simply enter the IP address of a Network Policy Server. You can skip this step if the Network Policy Server is running on the local machine.
9. **Click** Next.
10. The following screen asks for the NAP Enforcement Servers running DHCP Server. The DHCP servers are treated as RADIUS clients. Click the Add button.
11. Provide the names of your DHCP servers.
12. **Click** OK.
13. **Click** Next.
14. The following screen asks which scopes the policy should apply to. You can use this option if you want to apply different policies to different scopes. If you want the policy that you are creating to be applied to all NAP enabled DHCP scopes, you should leave the DHCP Scopes option blank.
15. **Click** Next.

16. The next screen asks you to provide any groups to which the policy should apply. Since the policy should apply to all users this screen should be left blank.
17. **Click** Next.
18. Specify a remediation server if you have one.
19. If you have a help URL, enter it into the space provided. Typically, the Help URL would point to a Web page explaining to users what they need to do to make their computer compliant with your policy.
20. **Click** Next.
21. Choose the NAP health policy that you want to use. By default, Windows includes a NAP health policy named Windows Security Health Validator.
22. **Select** the Enable Auto Remediation of Client Computers option.
23. **Choose** whether you want to allow or deny network access to incompatible network clients.
24. **Click** Next.
25. Verify that the information presented on the summary screen is correct.
26. **Click** Finish.

Configuring System Health Validators

The steps in the previous section created a NAP policy that granted or denied clients access to the network based on whether or not the clients met certain health criteria. The NAP policy was linked to a default system health validator named Windows Security Health Validator. You can use the steps below to configure the Windows Security Health Validator in a way that defines what it means for computers on your network to be healthy:

1. **Open** Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Network Policy and Access Services | NPS (Local) | Network Access Protection | System Health Validators.
3. **Double click** on the Windows Security Health Validator.
4. When Windows displays the Windows Security Health Validator Properties sheet, **click Configure**.
5. **Select** the check boxes for the health settings that you want to enable. For example, you can require a firewall to be enabled for all network connections and an antivirus application to be installed and up to date. It is worth noting that this properties sheet contains two separate tabs that maintain two separate sets of health settings. One group of settings applies to Windows Vista (and Windows 7) and the other applies to Windows XP (SP3 and higher).
6. When you have specified your desired health requirements, **click OK**.

Remote Access

From the very beginning, Windows Server has offered support for remote access. In the days of Windows NT, remote access was enabled through the use of dial-up networking. Windows Server 2008 still supports dial up networking through the Routing and Remote Access service; but, the preferred method for remote network access is to set up a VPN connection.

Windows Server 2008 can be configured to act as a VPN server. The server offers full support for PPTP (the Point to Point Tunneling Protocol) and L2TP (the Layer 2 Tunneling Protocol). This makes remote connectivity easy because Windows clients natively support both of these protocols.

One of the big drawbacks to the PPTP and the L2TP protocols is that they do not use typical firewall ports. As such, firewall ports must be opened in order to allow PPTP and L2TP connections. Even though there is nothing especially difficult about opening a firewall port, most networks are designed with multiple firewalls in place. As such, the necessary ports may need to be opened on several different firewalls.

Unfortunately, this problem is not limited solely to the corporate network. Many firewalls also block PPTP and L2TP traffic on outbound firewalls. For example, it is common for hotels to configure their Internet connections to block everything except for HTTP and HTTPS traffic.

Because of these challenges, Microsoft's preferred VPN protocol for Windows Server 2008 is a new protocol called SSTP (Secure Socket Tunneling Protocol). The SSTP protocol controls VPN traffic over a standard HTTPS sessions using port 443. This is similar to the technique that Microsoft uses for the TS Gateway in which they tunnel RDP traffic through an HTTPS session. Microsoft uses a similar technique in Exchange Server 2007 and Exchange Server 2010 to allow Outlook to connect to Exchange Server using an SSL encrypted HTTPS connection.

As was the case with PPTP and L2TP, the SSTP protocol is capable of using Extensible Authentication Protocol (EAP) authentication methods. However, before you can use SSTP for your VPN connections you must have a certificate infrastructure in place. SSTP requires SSL encryption, which in turn requires the use of a certificate.

It is important to remember that not all organizations use Windows Server 2008 as a VPN solution. Some organizations might use an older version of Windows Server or a non-Windows server, such as a Linux VPN. There are also hardware appliances that can serve as VPN gateways. Windows Server 2008 can be configured to interact with these types of VPNs by acting as a RADIUS server.

A RADIUS server is an authentication server that is used for VPN connections. When a user connects to the VPN server, the VPN server hands the user's credentials off to a RADIUS server for authentication. In Windows Server 2008 it is the Network Policy Server that acts as a RADIUS server. In the previous version of Windows Server RADIUS was offered through a service known as the Internet Authentication Service (IAS).

Configuring Routing and Remote Access

If you want to configure Windows Server 2008 to act as a VPN server, you will have to install and configure the Routing and Remote Access Services. The Routing and Remote Access Services are included as a part of the Network Policy Services role. The procedure below walks you through adding and configuring the Routing and Remote Access Services, but assumes that the Network Policy Services role is already installed:

1. **Open** the Server Manager.
2. Navigate through the console tree to Server Manager | Roles | Network Policy and Access Server.
3. **Click** the Add Role Services link.
4. **Select** the Remote Access Service check box.
5. **Click** Next.
6. **Click** Install.
7. **Click** Close.
8. Close and re-open Server Manager.
9. Navigate through the Server Manager's console tree to Server Manager | Roles | Network Policy and Access Services | Routing and Remote Access.

10. **Right click** on the Routing and Remote Access container.
11. **Choose** the Configure and Enable Routing and Remote Access command from the shortcut menu.
12. When the Routing and Remote Access Server Setup Wizard begins, **click Next** to bypass the Welcome screen.
13. **Choose** the Remote Access (Dial-Up or VPN) option.
14. **Click** Next.
15. **Select** the VPN check box.
16. **Click** Next.
17. When prompted, choose the network adapter that the server uses to connect to the Internet.
18. Make sure that the Enable Security on the Selected Interface By Setting up Static Packet Filters check box is selected.
19. **Click** Next.
20. When the wizard asks you how IP addresses should be assigned to clients, choose the Automatically option.
21. **Click** Next.
22. Tell the wizard whether or not you want to use RADIUS to authenticate VPN connection requests.
23. **Click** Next.
24. **Click** Finish.

Once Routing and Remote Access has been installed and enabled, there is still a bit of configuration work that needs to be done before it can be used. To complete the configuration process, follow these steps:

1. Navigate through the Server Manager console tree to Server Manager | Roles | Network Policy Access Services | Routing and Remote Access.
2. **Right click** on the Routing and Remote Access container and select the Properties command from the shortcut menu.
3. When the Routing and Remote Access Properties sheet opens, select the **IPv4** tab.
4. **Choose** to have name resolution performed by a DHCP server. As an alternative you can provide a static address pool. By doing so the Remote Access Server will act as a DHCP server for network clients. It is important, however, to make sure that the static IP address pool does not overlap any of the IP address scopes that have been defined on your DHCP server.
5. **Click** OK.
6. Navigate through the Server Manager's console tree to Server Manager | Roles | Network Policy and Access Services | Routing and Remote Access | Ports.
7. **Right click** on the Ports container and select the Properties command from the shortcut menu.
8. You will now see a screen showing you the types of connections that you are currently supporting as well as the total number of each type of connection that will be simultaneously supported. By default, PPPoE, PPTP, L2TP, and SSTP are all supported. If you are attempting to create a VPN link between two offices you will have to use PPTP or L2TP. However, if you are simply supporting user VPN connections, SSTP is the connection type of choice.
9. **Click** on a VPN connection type and click Configure.
10. Use the Configure screen to set the maximum number of connections that are allowed for that protocol. Remember that VPN servers running Windows Server 2008 Standard Edition only support a total of 250 connections.
11. **Click** OK.
12. **Click** Yes to continue.

Fault Tolerance

In any organization, there are certain applications that are considered to be mission-critical. It is important to provide fault tolerance for these applications so that the applications remain available to users even if a failure were to occur.

There are many different ways to achieve fault tolerance. There are both hardware and software solutions. From a Windows Server 2008 prospective, fault tolerance is achieved by providing high availability. Microsoft offers two main solutions for achieving high availability with Windows Server 2008. These solutions include Network Load Balancing (NLB) and Failover Clustering.

Network Load Balancing

Although NLB can be thought of as a high availability service, it is better suited to providing scalability for high demand applications such as websites. NLB is designed to split a workload between multiple servers within a NLB cluster. A NLB cluster can consist of up to 32 different servers, all running Windows Server 2008.

NLB does provide high availability for a clustered application through the use of redundancy. Each NLB cluster contains multiple servers; so if a server were to fail then another server can pick up the slack. However, there are some limitations to NLB that you need to be aware of.

The biggest limitation is that NLB does not provide data redundancy. It is only suitable for providing application redundancy. Imagine, for example, that you created a NLB cluster consisting of five different database servers. When a user decides to connect to your database application the NLB service will direct the user to one of those five servers. Now let's suppose that the user makes a change to the database. That change occurs locally on the server that the user has connected to; but, the change is not made to any of the other servers in the NLB cluster.

Because NLB does not replicate data between cluster nodes, the NLB service is only suitable for applications, not for data. In most cases, the NLB service either hosts a static application (such as a website) or each copy of the application is connected to a centralized database on a server that exists outside of the NLB cluster.

Although NLB does provide redundancy, it is usually treated more as a scalability solution. For example, imagine that you have a web front ends that is suffering from excessive traffic. NLB offers a perfect way to take care of the problem because the inbound workload can be divided up among all of the different servers within the NLB cluster. That way, no one single server has to carry the entire workload. If your workload increases over time, you can simply add additional servers to the NLB cluster as a way of dealing with the extra demand.

If you are going to use NLB as a scalability solution, you are going to need to do a little bit of planning before you build your NLB cluster. That's because you need to ensure that if one of your cluster nodes fails that the remaining servers have adequate resources to carry the load of the failed server. For example, if your web application receives so much traffic that it takes a minimum of two servers to handle all of the traffic, then you wouldn't want to create a two node NLB cluster. Otherwise, if one of the cluster nodes failed then the one remaining server would be inadequate for supporting the workload. In this type of situation, the NLB cluster should ideally be made up of three or four cluster nodes.

Failover Clustering

Another high availability solution offered by Windows Server 2008 is failover clustering. Failover clustering is different from NLB in many different ways. For starters, NLB is available in all editions of Windows Server 2008. In contrast, failover clustering is only available in the Enterprise and Data center editions. Another difference is that while NLB can support up to 32 nodes within a cluster, failover clustering has a limit of 16 nodes.

There are also major differences in what NLB and Failover Clustering are designed to do. As previously mentioned, NLB is designed primarily as a scalability solution. NLB clusters are designed to split a workload across multiple servers within the cluster. That way, the cluster can grow as the workload increases. NLB does achieve a degree of fault tolerance through redundancy, but is far from being a comprehensive fault tolerant solution. NLB provides fault tolerance and redundancy for applications, but not for data.

In contrast, failover clustering protects the data. You can think of failover clustering as a form of server redundancy. In other words, failover clusters have one or more spare server standing by. That way, if a server were to fail then one of the spare servers within the cluster (known as passive nodes) are able to take over automatically.

Another major difference between NLB and failover clustering is that failover clustering is not used for workload distribution. NLB is designed to evenly distribute a workload across multiple application servers. This is possible because none of the servers within the cluster contain a writable database (although they can be linked to a database that resides outside of the cluster). In contrast, nodes within a failover cluster can contain both applications and data. This makes workload distribution impossible.

The Majority Node Set

One of the most important concepts to understand with regard to failover clustering is that of a majority node set. The idea behind this concept is that occasionally it may be possible for the cluster to fail due to a network failure rather than the failure of an entire server. In this type of situation, all of the cluster nodes remain active; but, some of the cluster nodes are unable to communicate with other nodes. If left unchecked such a situation results in what is known as a split brain failure. A split brain failure means that a network failure has essentially divided the cluster into two separate but functional parts. Each remaining portion of the cluster thinks that the nodes that it cannot communicate with have failed. The end result would be two separate instances of the cluster.

To keep this from happening, Microsoft has designed failover clustering so that a cluster has to maintain quorum in order to function. This is where the term Majority Node Set comes into play. The basic idea is that the majority of the cluster nodes must still be functional in order for the cluster to achieve quorum. Microsoft defines the majority as half of the cluster nodes + 1. Therefore, if a failover cluster were made up of five cluster nodes, at least three of those nodes would have to remain online in order for the cluster to retain quorum.

So what would happen if a cluster has an even number of nodes? In that type of situation the half + 1 rule still applies. However, it is possible for a failure to occur in a way that prevents the cluster from retaining quorum. For example, if a network failure caused three nodes out of a six node cluster to become unavailable then you could potentially end up with two different clusters made up of three nodes each. Neither of the cluster segments would achieve quorum because the half + 1 rule would require four nodes to remain online. Therefore, a split brain failure would occur and the entire cluster would go down.

There are two ways to keep this from happening. One method is to avoid creating clusters with even numbers of nodes. The other method is to make use of a file share witness. A file share witness takes the place of a cluster node in clusters made up of an even number of nodes. If a split brain failure occurs, the cluster segments attempt to communicate with a designated file share on the network. Whichever cluster segment is able to access the file share, it is given quorum and continues to function.

Backup and Recovery

In any network, there will likely be times when data needs to be restored from backup. Windows Server 2008 provides tools for recovering from a disaster; but, the actual technique that you will use varies tremendously depending on what type of data it is that needs to be restored. For example, you might use one technique to restore a server from backup but a completely different technique to roll your Active Directory back to a previous state.

Active Directory Recovery

Although Windows Server 2008 does make it possible to restore a backup of the Active Directory, doing so is almost never necessary. Microsoft builds safeguards into the Active Directory that will usually allow you to avoid a full blown restoration.

One type of situation in which an Active Directory recovery might be necessary is in the event of a domain controller failure. However, as long as there are other domain controllers within the domain a domain controller failure is not usually catastrophic. You can simply restore the failed domain controller from backup. Of course, this results in the restored domain controller having an older version of the Active Directory database than what is currently on your other domain controllers; but, this isn't a problem. The other domain controllers simply replicate any missing Active Directory objects to the recently restored domain controller, bringing it up to date.

Another example of a situation in which an Active Directory recovery may be necessary is when an object (such as a user account) is accidentally deleted from the Active Directory. Thankfully, an accidentally deleted object can be recovered without having to roll the entire Active Directory database back to a previous state.

The reason why it is possible to recover a deleted object without restoring a backup is because when an Active Directory object is deleted, it is not actually removed from the Active Directory database. Instead, the object is tombstoned. This means that the Active Directory flags the object with a special attribute so that it will be treated as if it were deleted, but without actually being removed from the database. If you want to get the deleted object back, you perform a procedure called tombstone reanimation, which essentially removes the tombstone attribute from the object.

So if deleted objects are not actually deleted from the Active Directory database, what is to stop the database from growing until it eventually runs the server out of disk space? Well, tombstoned objects are eventually deleted. When an object is tombstoned, the Active Directory marks the date on which the tombstone was placed on the object. The domain controllers periodically run a process called garbage collection which permanently deletes objects that have been tombstoned for an excessive length of time.

By default an object remains tombstoned for 180 days. This time period (which is known as the tombstone lifetime) is the amount of time that you have to recover the deleted object before it is removed from the Active Directory database by the garbage collection process.

The Active Directory Snapshots and Mounting Tool

One of the big problems with tombstone reanimation is that when an object is tombstoned, the object itself remains in the Active Directory (flagged as a tombstoned object); but, most of the object's attributes are deleted. Therefore, if you ever decide to recover a deleted object from the Active Directory, you will have to do some Active Directory editing to fill in any attributes that have been stripped from the object during the tombstoning process.

Needless to say, this can be a big problem if you do not know what the attribute values are supposed to be. This is where the Active Directory Snapshots and Mounting Tool comes into play.

Windows Server 2008 uses a component called the Volume Shadow Copy Service to make snapshot backups of the Active Directory database (among other things). The Active Directory Snapshots and Mounting Tool allows you to mount an Active Directory snapshot backup. Once a snapshot has been mounted, you can't use the snapshot to recover the deleted object or its attributes. You can, however, use the snapshot to browse the Active Directory in an effort to track down the values that should be assigned to the recently recovered object's various attributes.

This tool can also be extremely helpful if data corruption should ever occur within the Active Directory. Data corruption is one of those situations that usually does warrant restoring the full Active Directory database.

If you find that you need to restore the Active Directory, time is of the essence. Everything on your network revolves around Active Directory; so, you will want to complete the restoration quickly. The problem is determining when the corruption occurred can be a guessing game. You can waste a lot of time if you restore a backup and then check to see if the corruption still exists. You may end up having to restore several backups before you find one that isn't corrupt.

The Active Directory Snapshots and Monitoring Tool gives you a way to quickly peek into your Active Directory snapshots so that you can determine when the corruption occurred before you restore a backup.

Recovering Encrypted Data

Windows Server 2008 supports the encryption of user data through the Encrypting File System (EFS). As such, end users can use EFS to encrypt any files or folders that they want to protect.

Although EFS provides for good security, it can result in data loss if you have an uncooperative user. When a user encrypts files and folders, Windows uses a set of keys that are tied to the user's account. This means that only the person who encrypted the files is able to decrypt them.

Because administrators sometimes need to access data that was encrypted by an end user, Microsoft has designed Windows Server 2008 to support the use of a Disaster Recovery Agent. A Disaster Recovery Agent is a user who has been given the authority to decrypt encrypted data for disaster recovery purposes. In smaller environments, the Disaster Recovery Agent is typically the domain administrator. In larger environments, the agent is often a dedicated security professional.

Disaster Recovery Agents can be designated through group policy settings. Before setting up a Disaster Recovery Agent, however, it is important to check with your corporate attorneys to make sure that having a Disaster Recovery Agent does not violate any compliance regulations.

BitLocker Encryption

BitLocker is designed to provide volume level encryption. Like EFS, BitLocker encryption is based on the use of keys. If these keys are lost, an encrypted volume will remain encrypted and inaccessible.

One way that you can protect your organization against the loss of BitLocker recovery keys is to back the keys up to Active Directory. Active Directory level backups of BitLocker recovery keys is not automatic. It must be enabled through a group policy setting.

Windows Server Backup

In Windows Server 2008, Microsoft has retired the NTBACKUP program that has been a part of Windows Server since the very beginning. NTBACKUP has been replaced by a new backup program called Windows Server Backup.

If you previously used NTBACKUP to back up Windows Server 2003 prior to upgrading to Windows Server 2008, then you need to know that Windows Server Backup will not restore a backup file (a .BKF file) that was created by NTBACKUP. Microsoft does, however, offer a free downloadable utility that will allow .BKF files to be restored to Windows Server 2008.

In order to create a backup using Windows Server Backup, you must be a member of either the Administrators group or the Backup Operators group. Remember that domain controllers share groups; so, if you add someone to one of these groups on a domain controller, they will be members of that group on every domain controller in the entire domain.

When you create a backup using Windows Server Backup, you are given several different choices of what you can back up. Your choices include the full system, OS volumes, or non-OS volumes. Windows Server Backup does not offer the ability to backup individual files and folders as was possible with NTBACKUP. Windows Server Backup is designed primarily for performing server level backups rather than granular backups.

If you create a backup using Windows Server Backup, the backup is written to a backup file as was the case with NTBACKUP. Windows Server Backup allows you to write this file to a local storage device, an external drive, or to a remote file share.

Windows Server Backup does not allow you to back data up to a tape drive. While the operating system itself does support tape drives, they are only supported for use with third party backup applications.

Installing Windows Server Backup

In previous versions of Windows Server, NTBACKUP was always installed by default. However, Windows Server 2008 does not install Windows Server Backup by default. Windows Server Backup is made available as an operating system feature. You can install Windows Server Backup by following these steps:

1. **Open** the Server Manager.
2. **Select** the Features container.
3. **Click** the Add Features link.
4. **Select** the Windows Server Backup Feature option from the list of available features.
5. If you expand the listing for the Windows Server Backup Feature, you will see that there is also a Command Line component that you can install. This component is useful for creating scripted backups, but does not need to be installed in order to achieve GUI backup capabilities.
6. **Click** Next.
7. **Click** Install.
8. When the installation process completes, **click Close**.

Creating a Scheduled Backup

After Windows Server Backup has been installed, it is a good idea to configure it to run a regularly scheduled backup. To do so, follow these steps:

1. **Select** the Windows Server Backup command from the server's Administrative Tools menu.
2. When Windows Server Backup opens, click the Backup Schedule link found in the Actions pane.
3. When the Backup Schedule Wizard launches, **click Next** to clear the wizard's Welcome screen.
4. **Choose** the type of backup that you want to create. You can either create a Full Server backup or a Custom backup.
5. **Click Next**.
6. If you are creating a Full Server backup then you will see a screen asking you how often you want to perform the backup and when.
7. After setting your backup schedule, **click Next**.
8. **Select** your backup destination.
9. **Click Next**.
10. Assign a label to the destination disk.
11. **Click Next**.
12. Verify that everything on the confirmation screen is correct.
13. **Click Finish**.

There are a few gotchas that you need to be aware of when setting up a backup with Windows Server Backup. First, if you are creating a scheduled backup, Windows will dedicate that entire disk to the backup process. You cannot use the disk for anything else. In fact, once the first backup has been created the chosen disk will even be hidden from Windows Explorer.

Another important thing to know about setting up a backup is that only NTFS volumes can be backed up. Windows Server Backup does not work with FAT or FAT-32 volumes.

Restoring Data

Windows Server Backup is based on the Volume Shadow Copy Service. Therefore, Windows Server Backup supports the restoration of individual files and folders as well as any VSS aware applications that might be installed on the server. Windows Server Backup also makes it possible to restore an entire volume or even the entire server.

The process of performing a full system restoration has been greatly improved in Windows Server 2008. In previous versions of Windows Server, if you wanted to perform a full system recovery, you had to install Windows onto the server, install the service pack that the previous operating system had been using, and then restore your data.

In contrast, Windows Server 2008 is capable of performing a true bare metal restore. There is no need to install Windows prior to restoring the backup. Simply boot the server from the Windows Server 2008 installation DVD and choose the option to recover your server. You then specify where the backup is located. You can point Windows toward an external hard drive or even to a remote share. Windows then uses the system image to perform a full bare metal recovery.

Troubleshooting

Every once in a while servers running Windows Server experience problems. When this happens, you need to know how to use Windows Server's built-in diagnostic tools to determine the source of the problem so that you can return the server to a fully functional state.

The Event Viewer

When a system crash occurs, one of the first things that you should check is the Event Viewer. Windows Server is designed to record as much information as possible into the System Log any time that a problem occurs. Often, you can use this information to determine the root cause of the failure.

The Performance and Reliability Monitor

Another tool that you can use to help troubleshoot problems is the Performance and Reliability Monitor. As you will recall, the Performance and Reliability Monitor offers numerous performance counters that you can use to track various aspects of the system's performance.

When it comes to looking for hardware problems, you can use the Performance and Reliability monitor to look for things like hard drives that are being used excessively even when the system is not under a load. You might also check for network broadcast storms. The Performance and Reliability Monitor can also be useful for troubleshooting software related problems too because many such problems are related to memory leaks (which can be tracked with this tool).

The Device Manager

One of the most beneficial tools for troubleshooting hardware problems is the Device Manager. The Device Manager is responsible for managing every hardware component in the entire system. As such, the Device Manager is able to report which devices are working and which are not.

The Device Manager also contains a mechanism for rolling back device drivers. This is important because sometimes Microsoft deploys updated hardware device drivers through Windows Update. On occasion these drivers have been known to overwrite manufacturer's drivers and cause a device to stop working. In these types of situations, the Device Manager can fix the problem by allowing you to roll back the device driver.

The Memory Diagnostic Tool

Sometimes a system's behavior seems to completely defy logic. When these types of really strange problems occur they can almost always be traced to memory issues.

Microsoft includes a utility that can be used to test the server's memory to be sure that it is reliable. This tool, which is called the Memory Diagnostic Tool, is located on the Windows Server installation DVD. You can access the tool by booting off of the Windows Server installation DVD and choosing the option to recover your server.

The Memory Diagnostic Tool is also installed with Windows Server by default. Although it is possible to initiate the tool from inside of the server operating system, the tool can't actually run until your system is rebooted. This is because the Memory Diagnostic Tool needs to be able to test memory that is normally in use by Windows. You can launch the Memory Diagnostic Tool by following these steps:

1. Choose the Memory Diagnostic Tool command from the server's Administrative Tools menu.
2. You will now see a prompt asking you if you want to restart the server now or if you want to schedule the tool to run the next time that you reboot the server.
3. Make your choice and then reboot the server. The Memory Diagnostic Tool will run as a part of the server reboot.

Security Troubleshooting

Occasionally, a system's behavior may lead you to wonder if perhaps a security breach has occurred. In these types of situations, you can verify your server's security by running Microsoft's Baseline Security Analyzer.

The Microsoft Baseline Security Analyzer works by comparing your server's configuration against Microsoft's best practices for security. The tool looks for missing or weak policy settings as well as missing security patches.

Practice Questions

Chapter 1

1. You plan to deploy 10 Windows Server 2008 servers and 75 Windows Vista workstations to a secure remote facility that is a part of your organization. This secure remote facility is not connected to the Internet. A low-speed WAN connection will link the headquarters with the secure remote facility. You plan to deploy operating systems in the remote facility by using Windows Deployment Services (WDS). However, you need to determine the appropriate Windows Activation model for your infrastructure. Network traffic over the WAN must be strictly limited.

Which of the following actions should you perform?

- A. Use RIS instead of WDS as an operating system deployment method.
 - B. Implement a VA 1.0 volume license key infrastructure.
 - C. Implement a KMS activation infrastructure for the secure remote facility.
 - D. Implement a MAK activation infrastructure for the secure remote facility.
2. Your company's Windows Server 2008 Active Directory domain is organized physically as a headquarters site and two branch office sites. The headquarters site contains three domain controllers, and each branch office site contains two domain controllers. You need to deploy DNS to ensure that host name resolution for all hosts, forest-wide and in both sites, is supported with the least amount of administrative effort. DNS fault tolerance is also a must.

Which of the following actions should you perform?

- A. Deploy the GlobalNames zone on an authoritative DNS server.
 - B. Create a standard primary zone on a DNS server in the headquarters site and a standard secondary zone on a DNS server in each branch office site.
 - C. Ensure that all domain controllers have the DNS server role installed. Create an Active Directory-integrated zone and replicate zone data to all domain controllers in both sites.
 - D. Create a standard primary zone on a DNS server in the headquarters site. Create an Active Directory-integrated stub zone on a DNS server in each branch office site.
3. You purchased one new 64-bit server computer that will replace two older 64-bit servers that host two 64-bit line-of-business (LOB) applications. You need to recommend an appropriate Windows Server 2008 edition and server role to support the new infrastructure.

Which of the following actions should you perform? (Select two choices. Each correct answer represents part of a single solution.)

- A. Install Windows Server 2008 Web Edition on the new server.
- B. Install Windows Server 2008 Standard Edition on the new server.
- C. Deploy Windows System Resource Manager (WSRM) on the new server.
- D. Deploy the Hyper-V server role on the new server.

4. Your organization consists of a single Windows Server 2008 Active Directory forest that contains seven domains arranged in a single tree. The root domain is named preplogicc corp.com. IT departmental mandate specifies that all domain controllers must have the DNS server role installed. You need to plan the DNS infrastructure for a new child domain named tech.east.preplogicc corp.com. The three chief requirements are that (a) DNS services are fault-tolerant in the tech.east.preplogicc corp.com domain, (b) the tech.east.preplogicc corp.com domain is automatically made aware of any new DNS servers that are brought online in the root domain, and (c) full forest-wide name resolution is supported.

Which of the following actions should you perform? (Select two choices. Each correct answer represents a part of a single solution.)

- A. Create a standard secondary zone for preplogicc corp.com.
- B. Create an Active Directory-integrated stub zone for preplogicc corp.com.
- C. Create a standard primary zone on tech.east.preplogicc corp.com. Next, configure conditional forwarding on this server for the remaining DNS zones in the forest.
- D. Create an Active Directory-integrated zone on tech.east.preplogicc corp.com.

Chapter 2

1. You manage a Windows Server 2008 Active Directory forest for your company. The forest consists of one domain that is organized into four sites with three domain controllers residing at each site. You need to delegate Group Policy Object (GPO) creation to members of the GPO_Dev domain global group. Members of this group need to be able to create and edit their own GPOs in the domain, but should not have the ability to link GPOs anywhere in the forest. (You will undertake this task yourself.) The GPO_Dev group should not be granted any unnecessary administrative privileges.

Which of the following actions should you perform?

- A. Add the members of the GPO_Dev group to the CREATOR OWNER group in Active Directory.
- B. Edit the NTFS permissions of the SYSVOL directory.
- C. Add the members of the GPO_Dev group to the Server Operators domain local group in Active Directory.
- D. Add the members of the GPO_Dev group to the Group Policy Creator Owners global group in Active Directory.

2. You manage a Windows Server 2008 Active Directory domain for a government research institute. In Active Directory Users and Computers, you have placed the user accounts for 10 Marketing department employees into an OU named MKTING. Inside the MKTING OU, you have also created a domain global group named Marketing_Team that contains the 10 employee user accounts. You need to grant the 10 employees the ability to (a) share folders, and (b) manage print queues on a departmental member server named THANATOS. Your solution must involve the least amount of administrative effort.

Which of the following actions should you perform? (Select two choices. Each correct answer represents a part of a single solution.)

- A. Create a domain global group in Active Directory and populate it with the 10 Marketing team user accounts.
 - B. Create a domain local group on THANATOS and populate it with the 10 Marketing team user accounts.
 - C. Add the MKTING OU to the local Administrators group on THANATOS.
 - D. Add the security group you created to the local Administrators group on THANATOS.
3. You manage your company's single Windows Server 2008 Active Directory domain. Your company establishes a branch office in a remote location where no domain administrators will be physically present. You plan to ship a server computer and installation media to the branch office and allow a tech-savvy employee who works at the branch office to (a) install Windows Server 2008 on the computer, and (b) join the computer to the corporate domain as an RODC. You must ensure that this individual has no additional privileges to other domain resources.

Which of the following actions should you perform? (Select two choices. Each correct answer represents part of a single solution.)

- A. Prestage a local administrator account for the branch office employee.
 - B. Prestage an RODC computer account in Active Directory.
 - C. Configure the RODC account in Active Directory to cache only the branch office employee's account credentials.
 - D. Delegate RODC installation privileges to the branch office employee.
4. You manage a single Windows Server 2008 domain that includes an IIS 7.0-based Web application that is exposed to your corporate intranet. You need to hire a contractor to perform maintenance on the site. The contractor will use a VPN connection to reach the Web server. However, for security reasons you do not (a) plan to create a domain user account for the contractor, or (b) plan to have the contractor connect to the Web server by using a Windows credential.

Which of the following actions should you perform to allow the contractor to remotely work on the Web site yet preserve security?

- A. Modify the .NET trust level of the specified Web application.
- B. Configure URL Authorization on the IIS Web server.
- C. Create an Application Pool for the contractor.
- D. Create an IIS Manager user account for the contractor.

5. You manage a Windows Server 2008 domain for a university research center. The forest and domain functional level is Windows Server 2008. All client workstations in the domain run Windows Vista SP1. The research associates in Dr. Walker's group have higher password security requirements than do the rest of the institute. Dr. Walker's group (user accounts and computer accounts) all reside in an organizational unit (OU) named WALKER in the domain. You must accommodate this special security situation by expending the least amount of administrative effort.

Which of the following actions should you perform?

- A. Define an Active Directory site for the WALKER OU.
 - B. Create a GPO that contains the Password must meet complexity requirements Group Policy setting and link it to the WALKER OU.
 - C. Define a new domain for Dr. Walker's research group.
 - D. Deploy a Passwords Settings Object (PSO) in the domain.
6. You manage a single Windows Server 2008 domain for your organization. You receive a support call from a user who complains that she is unable to access Control Panel items that she was able to open yesterday. Meanwhile, you remember that your colleague made changes to the domain's Group Policy infrastructure earlier today.

Which of the following actions should you perform to diagnose the user's issue?

- A. Use RSoP in Logging mode on the user's computer.
 - B. Use RSoP in Planning mode on the user's computer.
 - C. Define a WMI filter in the Group Policy Management Console on a domain controller.
 - D. Reset the user's computer account in Active Directory Users and Computers.
7. You manage a single Windows Server 2008 Active Directory domain for your company. You have created two GPOs, named GPO1 and GPO2, and linked them at the domain level. You do not want the settings for GPO2 to apply to the employees in the Human Resources department, whose user accounts are all stored in the HR organizational unit (OU). However, you need the Group Policy settings that are contained in GPO1 to apply to all domain users. Your solution must involve the least amount of administrative overhead.

Which of the following actions should you perform? (Select two choices. Each correct answer represents a part of a single solution.)

- A. Link GPO1 to the HR OU.
- B. Enable the Block Inheritance option on the HR OU.
- C. Unlink GPO2 from the domain.
- D. Select the Enforced option for GPO1.

Chapter 3

1. You manage a single Windows Server 2008 domain for a large health care management organization. You are designing a patch management strategy for your organization that uses Windows Server Update Services 3.0. The most critical business requirement of the plan is that server and client computers must be able to continue receiving updates in the event of a WSUS server failure.

Which of the following actions should you perform? (Select two choices. Each correct answer represents a part of a single solution.)

- A. Create a new Windows Internal Database (WID) failover cluster.
 - B. Create a new SQL Server 2008 failover cluster.
 - C. Configure WSUS for Network Load Balancing (NLB).
 - D. Configure WSUS for DNS round robin.
2. You manage a single Active Directory domain for a government contractor. All servers in the organization run Windows Server 2008, and all client computers run Windows Vista with Service Pack 1. IT security policy mandates that (a) all files must be stored in an encrypted state, and (b) files must be encrypted when they are transmitted over the LAN. You compose some schematic diagrams on your workstation. Next, you encrypt the files using Encrypting File System (EFS). Finally, you move the files to a file share on a member server named AJAX by using Windows Explorer.

Did a security breach occur? If so, what should you have done instead?

- A. A security breach did not occur. Your actions fell within your organization's IT security policy.
 - B. A security breach occurred. You should have applied EFS encryption to the files once they were moved to AJAX.
 - C. A security breach occurred. You should have secured the network media with IPsec.
 - D. A security breach occurred. You should have compressed the files before encrypting them on your local workstation.
3. You manage a single Windows Server 2008 domain for your company. Your organization consists of a main office and three branch offices. You set up a WSUS server in the main office that stores all Windows Update files locally. The branch offices are connected to the main office by a dedicated, high-speed WAN link. You want to configure a local WSUS server in each branch office. IT security policy mandates that all Windows Update approvals be handled centrally at the main office. Your solution must involve least administrative effort.

Which of the following actions should you perform in order to accomplish your goal?

- A. Create computer groups on each branch office WSUS server.
- B. Deploy each branch office WSUS server in offline mode.
- C. Deploy an autonomous mode WSUS server hierarchy.
- D. Deploy a replica mode WSUS server hierarchy.

4. You manage a single Active Directory domain for a private high school. All servers run Windows Server 2008 Standard Edition; all desktop computers and laptop computers run Windows Vista SP1. You want to be able to view and analyze the Event Logs for the school's six domain controllers from your administrative workstation in your office. It is important to see this data from a single interface. Another requirement is that the solution have minimal cost.

Which of the following actions should you perform to accomplish your goal? (Select two answers. Each correct choice represents a complete solution in itself.)

- Deploy WDS.
 - Create a blank MMC console. Load the Event Viewer snap-in for each of the school's domain controllers into the console.
 - Configure Event Forwarding.
 - Deploy Microsoft System Center Operations Manager (SCOM) 2007.
5. You manage a single Windows Server 2008 domain for a small university research institute. One of the institute's research teams works on a top secret government project, and therefore works from a disconnected network subnet. You have provisioned a WSUS server to this disconnected network subnet. You need to engineer a patch management strategy to keep this remote WSUS server current with updates arriving to the institute's primary WSUS server on the main network.

Which of the following actions should you perform in order to accomplish your goal? (Select two choices. Each correct answer represents part of a single solution.)

- A. Export the WSUS metadata from the SQL database.
- B. Export the XML metadata from IIS 7.0.
- C. Export the SQL database from the upstream WSUS server.
- D. Export the WSUSContent folder from the upstream WSUS server file system.

Chapter 4

1. You are part of an IT team that manages a large Active Directory domain that spans several geographic areas. All servers in the domain run Windows Server 2008. All client computers run Windows Vista. Your team has deployed a domain-based DFS root and configured replication to provide high availability for your users. After much deliberation, the team installed the primary DFS member in the headquarters site and one secondary DFS member in each branch office site, constituting a hub-and-spoke DFS topology.

Which of the following represents the chief disadvantage of this DFS topology?

- A. This DFS topology is more bandwidth-intensive than other DFS topologies.
- B. If the primary DFS member were to go offline, then replication would cease.
- C. Only one-way replication exists between the hub DFS server and the spoke DFS servers.
- D. Spoke replicas are not kept current with changes made on other spoke replicas.

2. You manage a single Windows Server 2008 domain for your organization, which is a medium-sized biotechnology company. Your company consists of a large sales force, all of whom carry laptop computers that run Windows Vista SP3. You need to design a secure Terminal Services remote access strategy for these users that supports the following criteria:

Users should be able to create a secure Remote Desktop connection to their desktop workstation at the headquarters office from anywhere in the world.

- All data between the remote user's computer and the corporate network should be encrypted.
- No additional port configuration/reconfiguration should be necessary at the corporate firewall.

Which of the following actions should you perform in order to accomplish your goal?

- A. Deploy a Windows Server 2008 member server with the TS RemoteApp role service on the network perimeter. Enforce policy with Microsoft ISA Server 2006.
 - B. Deploy a Windows Server 2008 member server with the TS Gateway role service on the network perimeter. Enforce policy with NPS.
 - C. Deploy a Windows Server 2008 member server with the TS Web Access role service on the network perimeter. Enforce policy with TS CAPs and TS RAPs.
 - D. Deploy a Windows Server 2008 member server with the TS Session Broker role service on the network perimeter. Enforce policy with WDS.
3. You manage a single Windows Server 2008 domain for a large auto parts distributorship. All servers run Windows Server 2008 Enterprise Edition. The client computer base represents a mix of desktop workstations and laptop computers that run Windows Vista Enterprise Edition with Service Pack 1. The infrastructure relies on Network Access Protection (NAP); a member server named NPS01 acts as a RADIUS server and Network Policy Server. You publish several line-of-business applications using TS RemoteApp on a TS Web Access server named WEB01. You need to provide secure remote access to these Terminal Services applications for a portion of your user population. You deploy a member server named ATHENA on the DMZ that has the TS Gateway role service installed, and you create a TS CAP and a TS RAP on ATHENA. You need to ensure that remote Terminal Services users are screened for compliance with NAP health policies.

Which of the following actions should you undertake?

- A. In TS RemoteApp Manager, configure the TS RemoteApp deployment settings with the option Do not use a TS Gateway server.
- B. In TS Web Access, change the Terminal Services server name to ATHENA.
- C. On NPS01, create a Connection Request Policy and specify WEB01 as the network access server.
- D. On NPS01, create a Network Policy and specify ATHENA as the network access server.

4. You manage a single Windows Server 2008 domain for a dry goods importer. All client computers in the domain run Windows Vista with Service Pack 1. Your development team has developed a custom inventory tracking application. In order to ensure centralized control over application updates, you decide to deploy this application via Terminal Services RemoteApp technology. To this end, you provision a Terminal Services member server named APP01 that hosts a TS Web Access site and publish the custom application via the TS Web Access interface. However, your users claim to have difficulty locating the TS Web Access site on the corporate intranet. Accordingly, you need to simplify the users' path to using the custom application.

Which of the following represent valid approaches to solving the problem? (Select two answers. Each correct choice represents a complete solution in itself.)

- A. Distribute the RemoteApp as an .RDP file to all users.
 - B. Distribute the RemoteApp as an .EXE installer to all users.
 - C. Publish the RemoteApp in a shared folder and notify the user population of the folder's location.
 - D. Publish the RemoteApp as a Start menu shortcut to all users.
5. You manage a single Windows Server 2008 domain for your company. The client base consists of Windows Vista and Windows XP SP2 computers. You deploy a member server named TS01 that is configured with the Terminal Services server role. TS01 hosts a TS Web Access Web site that offers a TS RemoteApp; the TS RemoteApp requires that Desktop Themes be enabled. You need to configure the network to support these requirements by expending the least amount of administrative effort.

Which of the following actions should you perform? (Choose two answers. Each correct choice represents part of a single solution.)

- A. Upgrade the Windows XP SP2 computers to SP3.
- B. Upgrade the Windows XP SP2 computers to Windows Vista.
- C. Install the Desktop Experience feature on TS01.
- D. Enable Themes on TS01.

Chapter 5

1. You manage a single Windows Server 2008 domain for a life insurance company. All client computers run Windows Vista. You perform a full server backup of a critical file server named XERXES every night. At 8:05 this morning, you discover that both the system and data volumes on XERXES have suffered catastrophic failures. You take XERXES offline and replace the drives with new standby disk drives. You now need to restore the operating system, all configuration settings, and (most importantly) all business data on XERXES in the shortest amount of time possible.

Which of the following actions should you perform in order to accomplish your goal?

- A. Use the Automated System Recovery (ASR) feature.
 - B. Use the Windows Recovery Console.
 - C. Use the Windows Server Backup tools.
 - D. Use the Complete PC Restore feature.
2. You manage a single Active Directory domain for your company. All servers run Windows Server 2008, and all client computers run Windows Vista. Your CTO asks you to deploy a high availability solution using Windows Server 2008 failover clustering. In particular, he asks you to define the difference between the technical terms failover and failback, which have historically confused him.

What should you tell your CTO? (Select two answers. Each correct choice represents a component of a single solution.)

- A. Failover occurs when the Windows Server 2008 failover clustering service moves and restarts an application's resources to the node that originally hosted the resource.
- B. Failover occurs when the Windows Server 2008 failover clustering service moves and restarts an application's resources from a failed cluster node to an available cluster node.
- C. Failback occurs when the Windows Server 2008 failover clustering service moves and restarts an application's resources from a failed cluster node to an available cluster node.
- D. Failback occurs when the Windows Server 2008 failover clustering service moves and restarts an application's resources to the node that originally hosted the resource.

3. You manage a single Active Directory domain for a private accounting firm. All client computers run Windows Vista. You plan to deploy an enterprise financial management application that uses Microsoft SQL Server 2008 for back-end data storage. You will install SQL Server on a Windows Server 2008 Enterprise Edition member server named PROTEUS. You need to design a disk storage strategy for SQL Server to (a) correspond with industry best practice, (b) provide for maximum application performance, and (c) provide data protection for the OS and the SQL data files in the event of a disk failure.

Which of the following actions should you perform in order to accomplish your goal?

- A. Store the OS and the SQL Server transaction logs on a RAID 0 array. Store the SQL Server data files on a RAID 1 array.
 - B. Store the OS and the SQL Server transaction logs on a RAID 1 array. Store the SQL Server data files on a RAID 5 array.
 - C. Store the OS and the SQL Server transaction logs on a RAID 1 array. Store the SQL Server data files on a RAID 0 array.
 - D. Store the OS and the SQL Server transaction logs on a RAID 1 array. Store the SQL Server data files on a RAID 1 array.
4. You manage a single Windows Server 2008 domain that consists of a main office and three branch offices. The WAN links that separate the offices are low speed and unreliable. Accordingly, you have designed an Active Directory physical topology in which each physical location is associated with an Active Directory site. All client computers run Windows Vista. You are planning to implement DFS in the domain. The two chief architectural requirements are as follows
- The DFS namespace must be fault tolerant.
 - Users must access only their nearest DFS namespace server.

Which of the following actions should you include in your implementation plan? (Select two answers. Each correct choice represents a part of a single solution.)

- A. Deploy a standalone DFS namespace.
- B. Deploy a domain-based DFS namespace.
- C. Set the referral option Exclude targets outside of the client's site for all folders in the DFS tree.
- D. Set the referral option Clients fail back to preferred targets for all folders in the DFS tree.

Answers & Explanations

Chapter 1

1. Answer: C

Incorrect A. In point of fact, WDS is the successor to the old and nasty Remote Installation Services (RIS) from Windows Server 2003/Windows 2000 Server.

Incorrect B. Windows Vista and Windows Server 2008 no longer use an untracked, use-as-many-times-as-you'd-like corporate volume license key anymore. Those days are gone. Welcome to Volume Activation 2.0!

Correct C. As long as you have at least 5 Windows Server 2008 computers and/or 25 Windows Vista machines in your environment, you qualify to host your own Key Management Service (KMS) activation server.

Incorrect D. Because the secure remote facility will contain more than 5 Windows Server 2008 computers and 25 Windows Vista computers, Microsoft requires that a KMS infrastructure be used instead.

2. Answer: C

Incorrect A. The GlobalNames zone is a new DNS zone type in Windows Server 2008 that is intended to make WINS largely irrelevant. However, the GlobalNames zone will not meet the requirements of this scenario.

Incorrect B. This "old school" solution provides little to no fault tolerance in the DNS infrastructure. Moreover, it is inelegant and difficult to administer, both initially and over time.

Correct C. By storing the DNS zone in Active Directory, you instantly guarantee fault tolerance as well as domain-wide and/or forest-wide host name resolution, thanks to AD replication.

Incorrect D. Stub zones are essentially "mini zones" that contain only enough resource records to "point" a DNS client to an authoritative server in another DNS domain. This solution only partially meets the requirements of the scenario.

3. Answer: B, D

Incorrect A. Windows Server 2008 Web Edition does not support the Hyper-V server role, which is required in this scenario. (That is, each LOB application will need to run as a virtual machine on the new server.)

Correct B. Windows Server 2008 Standard, Datacenter, and Enterprise Editions all support the Hyper-V virtualization server role. Note that Hyper-V requires a 64-bit (x64) processor architecture.

Incorrect C. WSRM is a server feature that enables an administrator to control the amount of processor and memory resources that a process or IIS application pool is allowed to consume on a single Windows Server 2008 computer.

Correct D. You should deploy Hyper-V on the new server and configure each LOB application to run inside of a separate virtual machine (VM) on that host server.

4. Answer: B, D

Incorrect A. Although this action would support forest-wide name resolution, it would not meet the requirement for fault tolerance.

Correct B. Stub zones contain only the name server (NS) records for the DNS zone in question. They also are dynamically updated when these records change.

Incorrect C. Standard DNS zones fail to meet the fault tolerance requirements as outlined in this scenario.

Correct D. By storing the domain's zone data in Active Directory, you achieve your requirement for fault tolerance. You also have the ability to replicate the zone data to all other domain controllers/DNS servers in the entire forest, providing forest-wide name resolution.

Chapter 2

1. Answer: D

Incorrect A. First of all, CREATOR OWNER is a special identity that does not appear in the groups list in Active Directory Users and Computers. Secondly, there exists an honest-to-goodness AD security group that fits your needs here: the Group Policy Creator Owners domain global group.

Incorrect B. Fortunately, you can delegate the proper permissions to the members of the GPO_Dev group in a much easier way than messing with the permissions of the SYSVOL public directory (which is almost never a good idea).

Incorrect C. Conferring this group membership upon the members of the GPO_Dev group would allow the users to log on to any domain controller locally and open the Group Policy Management Console. However, without membership in the Group Policy Creator Owners group, the users would still be unable to create and edit their own Group Policy Objects.

Correct D. By adding the appropriate user accounts to the Group Policy Creator Owners group, you delegate the ability to create and edit GPOs in the source domain to these users. However, the users cannot edit any other users' GPOs, nor can these users link their own GPOs to any domain object.

2. Answer: A, D

Correct A. A domain global group has forest-wide scope, although its membership is constrained to user accounts from the local domain.

Incorrect B. Because THANATOS is a member server and not a domain controller, you cannot create a domain local group, but only a local group. Local groups reside only on the local computer, and are therefore non-reusable elsewhere in the domain, making this option not optimal.

Incorrect C. This is impossible; you cannot, by design, add a domain-based organizational unit (or any OU, for that matter) to a group as a security principle.

Correct D. You should add the domain global group you created that contains the Marketing personnel to the local Administrators group on THANATOS. This action will immediately result in all Marketing personnel having full administrative rights on only that member server. The Marketing personnel will have no additional administrative privileges or user rights elsewhere in the domain.

3. Answer: B, D

Incorrect A. The user will need no credentials initially because he is performing a “bare metal” installation of Windows Server 2008 on new hardware. The user will, however, require his domain user account and a prestaged computer account defined in AD DS to complete the OS installation and the domain controller/RODC upgrade.

Correct B. If you plan to delegate the installation of an RODC, as is the case in this scenario, you will need to prestage an RODC account in Active Directory, and delegate the privileges to the appropriate domain user account(s); in this case, the branch office employee. This is a good example of the Administrator Role Separation (ARS) feature of Windows Server 2008.

Incorrect C. The Password Replication Policy is a security feature of Windows Server 2008 read-only domain controllers (RODCs) that reduces the exposure of the AD database to attack. However, this feature does not solve the problem of allowing the branch office employee to successfully join the computer to the domain and to promote the computer to become a domain controller.

Correct D. When you prestage the RODC computer account in Active Directory, you can grant the branch office employee’s domain account the privilege to install and/or locally manage the server without granting any additional domain-wide permissions.

4. Answer: D

Incorrect A. The .NET trust level for a Web application in IIS 7.0 deals with how much privilege a .NET-based Web application has on the local server and possibly the network, not how much privilege your contractor has on the server.

Incorrect B. Although URL authorization is a great feature in IIS 7.0, and you will want to use this feature to provide granular security to your content folder, this does not solve the problem outlined in this scenario; namely, how you can give the contractor remote administrative access to the Web application without a Windows domain account.

Incorrect C. An Application Pool in IIS is simply a memory construct (a container) that is used to protect the memory address space of a single, or more than one, Web application running on an IIS server.

Correct D. Internet Information Services (IIS) 7.0 supports internal user accounts known as IIS Manager user accounts that can be used for remote administrative connectivity to IIS servers and selected Web sites. These IIS Manager accounts can be allowed to view and/or use selected IIS features for specified Web application(s) running on the server.

5. Answer: D

Incorrect A. Active Directory sites are defined as physical locations that share high-speed network connectivity. Logically, AD sites contain only subnet objects and domain controller computer account objects, not OUs.

Incorrect B. Fine-grained password policies are defined by using ADSI Edit and the resulting Password Settings Container object in Active Directory Users and Computers. You cannot deploy multiple sets of password policies by including these settings in multiple traditional GPOs.

Incorrect C. With previous versions of Windows Server, this was the only method by which multiple password policies could be implemented in an AD forest. However, Windows Server 2008 introduces fine-grained password policies; this solves the problem nicely and reduces administrative burden to boot!

Correct D. You can define alternate password policies for specific users and/or groups by using ADSI Edit and the Password Settings Container in Active Directory Users and Computers.

6. Answer: A

Correct A. You should use Resultant Set of Policy (RSOP) in Logging mode by running Gpresult from an elevated command prompt on the user's computer in order to see precisely which Group Policy Objects are affecting the user and his or her computer.

Incorrect B. RSOP Planning mode, which is also known as Group Policy Modeling, is used for performing "what if" scenarios during the GPO design phase.

Incorrect C. You can run both "flavors" of RSOP from the Group Policy Management Console (GPMC). However, Windows Management Instrumentation (WMI) filters have utterly nothing to do with the situation at hand.

Incorrect D. Resetting a computer account in Active Directory is sometimes necessary when the computer's account password becomes out of sync in AD; however, in this case you are fairly certain that the problem lies with the GPO application, not a mismatched computer password.

7. Answer: B, D

Incorrect A. You can link a single GPO to more than one object in Active Directory. However, this increases administrative complexity and negatively impacts user logon speed and overall AD performance. Moreover, the Enforced option presents a better way to handle the GPO1 policy situation as described in this item.

Correct B. If you block policy inheritance on the HR OU, then all GPOs that are linked to upper levels in AD (in this case, both GPO1 and GPO2) will not be applied to the Human Resources employees. The scenario stipulates that the HR OU must not receive the GPO2 settings, and in this case, it will not.

Incorrect C. This is not a desirable option, because although you do not want the GPO2 policy settings to apply to your Human Resources employees, you do indeed need for these settings to apply to the rest of your staff. By unlinking the GPO entirely, nobody will receive the policy.

Correct D. When a GPO is marked as Enforced, the policy is applied to all downlevel AD containers, even if Block Policy Inheritance is enabled for that OU.

Chapter 3

1. Answer: B, C

Incorrect A. Although WSUS 3.0 can use the free Windows Internal Database that is built into Windows Server 2008, you cannot use WID in a failover cluster.

Correct B. WSUS 3.0 relies on a SQL Server database in order to function. To provide redundancy, you should create a failover cluster by using a full version of the product: either Microsoft SQL Server 2005 Enterprise Edition or Microsoft SQL Server 2008 Enterprise Edition.

Correct C. WSUS supports NLB, in which multiple WSUS front-end servers are load-balanced for reliability and performance reasons. If one WSUS server were to fail, other partner hosts continue operations. Because the SQL Server back-end is also in a cluster arrangement, the failure of a SQL Server cluster node does not bring down the database.

Incorrect D. Although Domain Name System (DNS) round robin works hand-in-hand with Network Load Balancing (NLB), NLB requires additional work, such as configuring the WSUS front-end servers all to point to the SQL Server cluster address.

2. Answer: C

Incorrect A. Encrypting File System (EFS) does not remain in effect over the network. Therefore, while the document bits travelled over the network medium, a man in the middle with a packet sniffer had plaintext access to the files. A local solution could be using IPSec; a remote access solution could be Secure Socket Tunneling Protocol (SSTP).

Incorrect B. Again, the movement of the files from your local workstation to the file server is the weakness here; the encryption goes away during that time of network transmission.

Correct C. Here you finally recognize the vulnerability inherent in EFS encryption and implement a workable solution, namely, IPSec policy on the LAN.

Incorrect D. An old truism with EFS encryption and NTFS compression is that these advanced NTFS attributes are mutually exclusive. That is to say, a file or folder on an NTFS volume can be encrypted or compressed, but never both together.

3. Answer: D

Incorrect A. This is precisely what you do not want; you need to centralize computer group and Windows Update approval status. Therefore, you need to deploy WSUS replica servers in the branch office locations.

Incorrect B. In high-security environments, you can set up a WSUS server in a disconnected network, export updates from its upstream neighbor, and then manually update the disconnected downstream neighbor. However, this method is not workable in this situation, nor is it necessary.

Incorrect C. In an autonomous mode WSUS hierarchy, the downstream WSUS servers (the branch office servers) would share the update packages with the main office WSUS servers, but the branch office admins would have full control over their own approvals, computer groups, and so on.

Correct D. A WSUS replica shares the update package database with its upstream partner and also shares all of the upstream partner's computer group and approval information. This topology is good for centralized IT departments, such as the case in this scenario.

4. Answer: B, C

Incorrect A. Windows Deployment Services (WDS) is free; it is bundled with Windows Server 2008. However, WDS is used to automate the deployment of operating systems, and has nothing whatsoever to do with the goals outlined in this item.

Correct B. A single Microsoft Management Console (MMC) shell can host many snap-ins, even several instances of the same snap-in.

Correct C. By configuring Event Forwarding and Event Subscriptions, you can receive some or all Event Log data from all domain controllers on your administrative workstations. By default, the Event Log data will appear in the Forwarded Events folder in the Event Viewer MMC console on your administrative workstation.

Incorrect D. SCOM 2007 is the successor to Microsoft Operations Manager (MOM) 2005. Both of these tools offer enterprise-level Event Log management. However, they are extremely expensive and this scenario calls for a low- or no-cost solution.

5. Answer: A, D

Correct A. You can use the `wsusutil.exe` command-line tool to take a backup of the WSUS metadata information from the upstream WSUS server in preparation to take and restore on the disconnected WSUS server.

Incorrect B. WSS 3.0 SP1 is the first version of Windows Server Update Services not to rely on a Web interface and Internet Information Services (IIS)-this is a good thing. Therefore, when exporting data to update an offline WSUS server, you do not need to monkey around with IIS at all.

Incorrect C. All you have to do is (a) export the metadata by using the `wsusutil.exe` (this tool queries the SQL database for you), and (b) export the update packages from the file system to prepare your WSUS updates for the disconnected downstream WSUS server.

Correct D. You can use Xcopy, the good old-fashioned Ntbackup utility, or any file-copy utility; the bottom line is that you need the contents of the `WSUSContent` folder to nab all of the Windows Update packages to ship to the disconnected WSUS server.

Chapter 4

1. Answer: B

Incorrect A. The hub-and-spoke Distributed File System (DFS) topology is far less bandwidth-hungry than the other common DFS, topology, the full-mesh topology.

Correct B. The “fatal flaw” with the hub-and-spoke DFS replication model is its single point of replication failure. On the other hand, hub-and-spoke is very efficient with regard to bandwidth because spoke servers replicate only with the hub server.

Incorrect C. Bidirectional replication takes place between the hub DFS server and the spoke DFS servers in the hub-and-spoke DFS topology.

Incorrect D. This is utterly untrue because bidirectional replication takes place between each spoke server and the hub server. Although the spoke DFS servers do not communicate directly, updates always flow through the primary hub server, and convergence eventually occurs for all servers in the entire DFS topology.

2. Answer: B

Incorrect A. The TS RemoteApp role service enables administrators to publish applications via Windows Server 2008 Terminal Services. TS RemoteApp is not a secure method for conducting Remote Desktop Protocol (RDP) sessions from outside the private internal network. However, you can indeed enforce policy with TS Gateway by using ISA Server 2006.

Correct B. The Windows Server 2008 TS Gateway server role uses standard Web ports (TCP 443, Secure Sockets Layer [SSL]) to support encrypted RDP sessions from anywhere in the world into a protected private intranet. TS Gateway can be used in conjunction with Microsoft ISA Server or Network Access Protection (NAP) for policy enforcement.

Incorrect C. TS Web Access is simply a browser-based interface for accessing Terminal Service applications that have been deployed by using the TS RemoteApp framework. You need to deploy the TS Gateway role service in this scenario. You can enforce TS Gateway policy by using Network Policy Server with TS CAPs and TS RAPs.

Incorrect D. The TS Session Broker role service is used to load balance a multi-server Terminal Services farm. If, however, you deploy TS Gateway to support secure over-the-Internet RDP connections, you can never enforce connection and authorization policies by using Windows Deployment Services (WDS); this is sheer nonsense.

3. Answer: D

Incorrect A. This is a trick answer choice. In contrast, you need to make sure that you have the role service communicating with TS Gateway in your TS RemoteApp Manager properties.

Incorrect B. You do not need to change the name of the TS Web Access Web server. Your remote access users will be able to establish connections to their TS RemoteApps from outside the private internal network in the very same manner as if they were sitting at their desks within the corporate campus.

Incorrect C. In this case you definitely want to ensure that a Connection Request Policy exists that specifies that connection requests sourced from ATHENA (not WEB01, which is the TS Web Access server) are forwarded to the Network Policy Server for screening against your NAP health policies.

Correct D. On the TS Gateway server, you will use TS CAPs and TS RAPs to perform initial connection authentication and resource authorization. Next, the Network Policy Server will screen the user and/or computer account against any configured NAP health policies.

4. Answer: A, D

Correct A. TS RemoteApps are extremely versatile. Remember that TS RemoteApps are never installed on the users' computers; therefore, when a user double-clicks the .RDP file, the application opens from the Terminal Services server, not on the local computer. However, the application looks and behaves as if it were installed locally on the user's machine.

Incorrect B. When you create an MSI package of a TS RemoteApp, all you are doing is making it easy for a user (or Group Policy Software installation, or SCCM 2007, as the case may be) to install a local .RDP file on the user's computer. The point with TS RemoteApp technology is that a "thick" application is never resident on the user's computer.

Incorrect C. This answer choice is not as laughable as you might think at first blush. You can indeed populate a shared folder with .RDP files or .MSI installer packages and direct your users to "nab" these files and copy them to their own computers. The .RDP files will then enable the users to launch the RemoteApps from their desktops.

Correct D. You can deploy TS RemoteApps in a variety of different ways. For instance, you can create a Windows Installer .MSI package and publish the RemoteApp as a Start menu shortcut. Moreover, the RemoteApp can be configured to start whenever the user double-clicks a file with an extension that is associated with the RemoteApp.

5. Answer: A, C

Correct A. Upgrading the Windows XP Service Pack 2 client computers to Service Pack 3 will give them the Remote Desktop Connection 6.0 Terminal Services client, which is fully interoperable with Windows Server 2008 Terminal Services.

Incorrect B. This solution, although it would certainly work, requires additional cost and far more administrative effort than is needed.

Correct C. In order to enable Desktop Themes for TS RemoteApps, you must install the Desktop Experience feature through Server Manager or by using the Servermanagercmd.exe command-line tool on the Terminal Services server.

Incorrect D. This answer choice points to a "chicken or the egg" issue: you cannot enable Desktop Themes on TS01 unless and until you install the Desktop Experience feature on the Terminal Services server.

Chapter 5

1. Answer: D

Incorrect A. ASR was a really ill-conceived feature in Windows Server 2003 that attempted to make it easier to recover the operating system (not user data) in the event of a catastrophic boot disk failure. ASR was replaced with the Complete PC Backup and Restore feature in Windows Vista and Windows Server 2008.

Incorrect B. Again, this answer choice represents a “red herring” intended to trip up those who are familiar with Windows Server 2003. The Recovery Console is completely gone from Windows Server 2008; instead, you use the product DVD to boot into Windows Recovery Environment (Win RE), a graphical, 32-bit mode in which you have access to the mouse and potentially the network as well.

Incorrect C. The point of this scenario is that the local operating system is destroyed. Therefore, you have no local access to any of your OS tools. Instead, you must boot the computer by using the Windows Server 2008 installation media and access the recovery tools available to you in the Windows Recovery Environment (Win RE).

Correct D. In Windows Server 2008 and Windows Vista, you can boot the computer by using the product DVD and initiate the Complete PC Restore feature. So long as you have a valid full server backup, you can restore the entire server to the state it was in as of the backup date. The full-server backup option uses disk-imaging technology not all that different from that of Symantec Ghost.

2. Answer: B, D

Incorrect A. This is the correct explanation for failback, not failover.

Correct B. Windows Server 2008 failover clustering is a built-in server role that provides high availability for services and applications.

Incorrect C. This definition is the correct explanation for failover, not failback.

Correct D. Failback represents the restoration of service after a failover process has taken place and the failed node is brought back online.

3. Answer: B

Incorrect A. The database will run very slowly if all writes are written in parallel fashion with a RAID 1 mirror set instead of a RAID 5 array. Moreover, RAID 0 provides excellent read/write performance, but no fault tolerance.

Correct B. Because recoverability is key for the operating system and the database transaction log files, you want to place those files on a mirror set, also known as a RAID 1 array. Because read performance plus fault tolerance is key for the database, RAID 5 makes sense for the SQL Server data files.

Incorrect C. Here you are okay placing the OS and the T-logs on a mirror set. However, storing the data files on a simple stripe set (RAID 0) will give you great performance for the database but absolutely no fault tolerance. This is a deal breaker as far as the scenario requirements are concerned.

Incorrect D. Although the OS and T-logs work just fine on a RAID 1 mirror set, the SQL data files will experience much faster write and read performance if stored on a RAID 5 stripe set due to parallel disk writes and reads.

4. Answer: B, C

Incorrect A. A standalone DFS namespace root is available only on the server in which it is created. Although you can replicate individual folders in the tree, if the primary server goes down, the entire DFS namespace is unavailable.

Correct B. Perhaps the biggest “selling point” of domain-based DFS is fault tolerance. You can create a fully redundant DFS root namespace, such that if one DFS root server goes down, the DFS root namespace (\\domainname\public, for instance) continues to operate as if no outage occurred.

Correct C. By configuring this referral option on all DFS links, a user will be restricted from accessing any DFS namespace servers outside of his or her Active Directory site.

Incorrect D. This option allows users to connect to DFS replicas in other AD sites. However, if the WAN link were to fail, the user would be transparently connected or reconnected to a preferred DNS server in the local AD site.