netwrix

BRIAN SVIDERGOL

MCITP, MCSE, RHEL3, VCP,
NCIE-SAN, MCT, MCSA,
Microsoft Certified
Solution Expert

70-410 EXAM STUDY GUIDE

Table of Contents

Introduction

This Study Guide covers Microsoft's 70-410 exam titled "Installing and Configuring Windows Server 2012" is the first exam of three exams that make up the Microsoft Certified Solutions Associate (MCSA) certification. The 70-411 exam and 70-412 exam are the other two exams. From there, you can continue taking tests to obtain your Microsoft Certified Solutions Expert (MCSE) certification, if desired. There are 9 MCSE certifications which each one being tied to a technology stack. You can take exams 70-413 and 70-414 to obtain the MCSE: Server Infrastructure. For server administrators, the MCSE: Server Infrastructure certification is a popular choice because the skills measured on the exam are closely tied to real-world job skills.

How this Study Guide is organized

The 70-410 exam is based on an exam object domain (OD). An OD is a blueprint for the exam. For 70-410, there are six sections. The sections are known as functional groups (FGs). Under each FG, there are subsections which are known as objectives. And finally, under objectives, there are lists of technologies or administrative tasks that the exam may test on. These are often called "may-includes". And the exam may or may not include them. They are meant to give test takers an idea of what topics are relevant and being tested but not an exact list of what is and what is not covered. More about that in the "How exam items are written" section. This Study Guide is organized like the exam OD. Every FG, objective, and "may include" is covered in this guide.

What this Study Guide is meant for

This Study Guide is not meant to be the only source of study material for the exam. It is not meant to be a guide that replaces in-depth books on topics or hands-on work in lab environments. It is not long enough to dive deep into a technology's history, its functionality, or its future. Instead, it is meant as a supplement to other study material. And when combined with other study material, this guide should help put you at an advantage when you take the exam.

How to use this Study Guide

If we group test takers into either the hands-on experience category or the inexperienced category, it helps dictate a good approach to study for the exam, as follows.

- **Hands-on experience**. In this category, administrators have spent time in their work history or at their current job working with Windows Server 2012 R2 and many of the associated technologies covered in the 70-410 exam. In other cases, administrators have hands-on experience with previous versions of Windows Server but their hands-on time with Windows Server 2012 R2 is from a lab or personal environment. I recommend the following study approach for this group:

    - Watch the 70-410 Exam Prep session from Microsoft Ignite.
    - Read this Study Guide. Use a highlighter to mark areas that you feel uncomfortable in.
    - Perform the commands and tasks shown in this Study Guide in a lab environment.
    - Spend extra time with the areas that you marked with a highlighter by reviewing TechNet, Microsoft Virtual Academy (MVA), and TechNet Virtual Labs. Start by reading about the technology and then spend a few minutes performing relevant tasks in a lab.
    - Optionally, try taking the 70-410 practice exam that is authorized by Microsoft. If you fail, mark the areas that were the most trouble and spend additional time in your lab.
    - On the day of the exam, re-read this Study Guide but skip areas where you feel strongest.

- **Inexperienced**. In this category, administrators may work in desktop support or another IT department that does not manage Windows Server or related technologies. Often, administrators in this category are exposed to the technologies from the client side. In some cases, these administrators have experimented with Windows Server in their personal lab. I recommend the following study approach for this group:
    - Buy and read a book on Windows Server 2012 R2 or a book dedicated to the 70-410 exam. A book will dive into more detail than this Study Guide and allow you to soak up more information from this Study Guide.
    - Watch the 70-410 Exam Prep session from Microsoft Ignite.
    - Read this Study Guide. Use a highlighter to mark areas that you feel uncomfortable in.
    - Spend extra time with the areas that you marked with a highlighter by reviewing TechNet, Microsoft Virtual Academy (MVA), and TechNet Virtual Labs. Start by reading about the technology and then spend a few minutes performing relevant tasks in a lab.
    - Take the 70-410 practice exam that is authorized by Microsoft. If you fail, mark the areas that were the most trouble and spend additional time in your lab.
    - On the day of the exam, re-read this study guide but skip areas where you feel strongest.

Avoid brain dumps or other sources of illegal test preparation material. Remember, the goal is to actually learn the material as part of the exam preparation. If you bypass the learning process to pass the exam then you are hurting yourself in the long run by reducing your value in the marketplace. Ultimately, it is the skills that you provide that determine your value.

How Microsoft exam items are written

An entire book could be written about the Microsoft exam development process. Here, the goal is to give readers a very brief introduction to the item writing part in this list of ten things that you may not know about Microsoft exam development:

1. Microsoft pays independent IT experts to write exam items. They use people from all over the world to bring a real-world perspective to exams.
2. Item writers use the exam OD published on the exam's web site to write items to. Item writers are not given access to anything more than what is available to the public.
3. Microsoft certification exams follow ISO 17024 which mandates that training development and exam development be separated and independent. If you ever wondered why an exam preparation book didn't cover all of the items that you saw on a certification exam, it is because of the separation. SMEs that write books and training material for the 70-410 exam are not allowed to participate in the 70-410 exam development process. And, 70-410 exam item writers are not allowed to be involved with the creation of 70-410 training material.
4. Microsoft prohibits the use of tricky questions or answers. For example, you won't see a PowerShell command that is perfect except that the number 1 is replaced with a capital I or a command that is missing a character such as a period that is tough to spot. If you think you saw a tricky item, there is a good chance that you don't know enough about the technology.
5. All items written are reviewed multiple times before they make it to the live exam. The review process often involves multiple SMEs independently reviewing an item. And, as a final review, items are evaluated during a beta run of the exam (although not all exams have a beta). Beta exams help to weed out any items that got past the initial review processes.
6. Exams have a sustained engineering component. A good example of sustained engineering is when the 70-410 exam was updated to reflect changes in Windows Server 2012 R2. During sustained engineering, poorly performing items (those that aren't accurately measuring a test taker's skills) are removed from the exam and new items (especially for product updates) are introduced into the exam item pool.
7. All of your comments and feedback that you enter when you take the exam are read by Microsoft. And they use those comments as part of the sustained engineering efforts.
8. You can participate in the exam development process if your skills in the exam technologies are at the expert level. You just need to fill out an SME profile form on the Microsoft Connect site at http://connect.microsoft.com/site862. Thereafter, you may get contacted for exam development opportunities.

1 - Install and configure servers (17%)

This functional group covers core installation and configuration topics centered around Windows Server. In addition, there is a section dedicated to storage technologies. This functional group is not tied to specific roles or services. Most administrators will have some experience with bits and pieces of each section. Plan to spend the most time on topics that you don't have experience with.

Install servers

This section, which did not change when changes were made to the exam OD for Windows Server 2012 R2, focuses on installation of servers, upgrades, and role migrations.

Plan for a server installation

This exam topic covers the steps you take prior to installing Windows Server. You should be comfortable with the different versions of Windows Server as well as the prerequisites. Below are key points to know for the exam:

- **Know the primary editions of Windows Server 2012 R2: Standard and Datacenter**. They both have the same hardware limits and support the same roles and features. There are two key differences:
    - **Automatic Virtual Machine Activation (AVMA)**. AVMA is only available on a virtualization host that runs Windows Server 2012 R2 Datacenter edition. It automatically activates VMs even if they don't have an internet connection.
    - **Virtualization rights**. When you use the Standard edition of Windows Server 2012 R2 on a virtualization host, you are able to use the Windows license to run 2 VMs. When you use the Datacenter edition on a virtualization host, you are able to use the license to run an unlimited number of VMs. Things can get a bit more complicated when factoring in multiple processors and VM movement between hosts. For more information see the Licensing Windows Server 2012 R2 for use with virtualization technologies white paper. Note that the information presented in this Study Guide likely covers what you need to know for the exam so don't spend too much time studying the white paper for the exam!
- **Know the minimum hardware requirements for Windows Server**. Note that all editions of Windows Server have the same requirements:
    - **1.4 GHz 64-bit processor**. This is the minimum processor supported.
    - **512 MB of RAM**. This is the minimum amount of RAM needed.
    - **32 GB of disk space**. This is the documented minimum. You can use less but it would be unsupported.

- **Also know some small facts about the planning process**. Know the following:
  - If you plan to perform a default installation of Windows Server 2012 R2, you will end up with a Server Core installation.
  - If you are installing VMs, you may need to allocate more than 512 MB of RAM for the installation or use dynamic memory. After installation, you can use 512 MB of RAM.

Plan for server roles

Prior to Windows Server 2012, planning for roles required quite a bit of knowledge around which roles were supported on which version of Windows Server. However, since Windows Server 2012, all roles are supported on the Standard edition and the Datacenter edition. So for this section, you should focus on the knowledge needed to install roles and manage roles.

Use the Install-WindowsFeature cmdlet to add roles and features. You must use the correct role and feature name which is often different than the one you see in Server Manager. Run the Get-WindowsFeature command to list all of the roles and features along with the name. Then, you can use the name to install a feature. For example, you can run the Add-WindowsFeature -Name Telnet-Client to add the Telnet Client feature. There are two parameters that you should know about:

- **The -IncludeAllSubFeature parameter**. This parameter will add the role and include all of the role services that are available under the role. While this parameter will install the management tools, you should not use it if you only want the management tools, because you will get more than what you need.
- **The -IncludeManagementTools parameter**. This parameter will add the associated management tools along with the role. If you install a role and can't find the management tools thereafter, it is likely because they were not installed along with the role.

You should also know how to view the installed roles and features. You can do that by looking at Server Manager. Or, you can use PowerShell. To view the name of all roles and features currently installed, run the **Get-WindowsFeature | where {$_.Installed -eq $True} | select Name** command.

Plan for a server upgrade

A server upgrade is a very precise process. It involves updating the existing operating system, such as Windows Server 2012, to a newer version such as Windows Server 2012 R2. Existing applications, services, and the configuration remain the same. Existing data remains as is. Administrators often refer to an upgrade as an in-place upgrade. There are some key limitations that you need to know about for the exam:

- **You can't upgrade a 32-bit Windows installation to a 64-bit Windows installation**. To move from a 32-bit Windows installation to a 64-bit Windows installation, you have to perform a new installation or perform a migration. In both cases, existing applications, services, and the configuration have to be reconfigured from the ground up.

- **You can't change languages as part of an upgrade**. If you have a US English language installation of Windows Server 2012, you cannot upgrade to the French language while upgrading to Windows Server 2012 R2.
- **You cannot downgrade the operating system edition during an upgrade**. If you have Windows Server 2012 Datacenter, you cannot upgrade to Windows Server 2012 R2 Standard.
- **You can upgrade the operating system edition during an upgrade**. If you have Windows Server 2012 Standard, you can upgrade to Windows Server 2012 R2 Standard or Datacenter.
- **To upgrade from Windows Server 2008 R2 to Windows Server 2012 R2, you need to have SP1**. You must have SP1 installed on Windows Server 2008 R2 prior to the upgrade.

Install Server Core

The Server Core installation is the default installation. The system requirements are almost identical. The exception is disk space because Server Core consumes about 4GB less disk space than a full installation.

Optimize resource utilization by using Features on Demand

Features on Demand is a technology introduced in Windows 8 and Windows Server 2012. It enables administrators to add roles and features without needing the installation media. A data repository containing the needed files for adding roles and features, named the component store, is located at %systemroot%\WinSxS. You can remove some or all of those files. The benefits of doing so are:

- **Reduce the amount of disk space used**. The side by side store can take up about 1.5 GB of space. When viewed in File Explorer, it appears to take up over 5 GB of space but that is misleading due to hard links that are in use. You can reduce the amount of disk space used, which can be helpful in situations where there is not ample disk space on a server. In the real world, this isn't much of a concern because disk space is very inexpensive and most servers have more than needed. Or, with virtualization, you can add disk space on demand.
- **Improve security (slightly)**. By removing the files, you make it more difficult to add roles and features to the server. In addition, you reduce the amount of files and code on the computer.

You can remove the files on a per role or feature basis or remove all of them. For example, to remove the Telnet-Client feature's files from a server, you can run the **Remove-WindowsFeature -Name Telnet-Client -Remove** command.

To get a detailed breakdown of the actual disk space being used by Features on Demand, go to a command prompt and run the **Dism /online /cleanup-image /analyzecomponentstore** command.

I expect this particular topic to go away in future revisions of the exam because it doesn't warrant a dedicated topic, especially with savings of about 1.5 GB of disk space. However, because it is in the OD, exam questions were written for it. Technically, those items may no longer be in the live exam since there are situations where items are invalidated (too hard, too easy, etc.).

Migrate roles from previous versions of Windows Server

The primary focus of this topic is the migration tools that Microsoft offers to migrate roles between servers, especially from server running an older version of Windows Server. The Windows Server Migration Tools feature provides PowerShell cmdlets to migrate roles from previous versions of Windows Server to Windows Server 2012 or Server 2012 R2.

The high level steps to migrate a role from Windows Server 2008 R2 or Windows Server 2012 to Windows Server 2012 R2 are:

1. Install Windows Server Migration Tools on destination server.
2. Create a deployment folder on the destination server. For example, to add a folder for a source server that runs a 64-bit installation of Windows Server 2012 in the C:\tempdeploy folder, you would run the **smigdeploy.exe /package /architecture amd64 /os ws12 /path c:\tempdeploy** command.
3. Copy the deployment folder to the source server.
4. Open an elevated command prompt on the source server, navigate to the deployment folder, and then register the source server by running the **smigdeploy.exe** command.
5. On the source server, perform role specific migration tasks. Depending on which role you are migrating, the tasks will vary. In general, the migration tasks involve copying data, migrating user accounts and configuration information. You use the Export-SmigServerSetting cmdlet to export data to the migration store.
6. On the destination server, use the Import-SmigServerSetting cmdlet to import configuration information.
7. Start services on the destination server. Migration complete.

Microsoft has put together extensive documentation on specific tasks for migrating roles. For the exam, do not memorize how to migrate all of the roles. Instead, read the high level steps for a few of them so you are familiar with the process. See https://technet.microsoft.com/en-us/library/dn486773.aspx for details on the specific role tasks.

There are a couple of other important items to know for the exam:

- **You may need to open the firewall for a role migration**. A role migration utilizes UDP and TCP port 7000.

- **You can migrate roles from older version of Windows Server**. The migration tools support a source server running Windows Server 2003 with SP2 and newer.
- **There are benefits of using the Windows Server Migration Tools over a manual migration**. The benefits are:
    - **Reduce risk by automating**. The less manual tasks an administrator has to perform, the less risk involved with migrations.
    - **Migrate roles faster**. The migration tools can help speed up the overall migration time which reduces downtime and impact.
    - **Migrations can be delegated**. Because the migration tools simplify the migration process, less experienced administrators can migrate roles.

Configure servers

This section is focuses on basic Windows Server configuration tasks. If you haven't used the Server Core installation, you should plan on spending some time with it as part of your exam preparation. When the exam OD was updated for Windows Server 2012 R2, a new topic titled "Install and configure Windows PowerShell Desired State Configuration (DSC)" was added.

Configure Server Core

The first time an administrator signs into a Server Core installation of Windows Server, they are often surprised at how sparse the interface is. The only thing to greet you is a command prompt. There isn't a desktop, a taskbar, or any menus. Configuring Server Core is different. For the exam, it is important that you know how to configure a Server Core installation, especially the initial setup. Here are important tips for the exam:

- **Know how to use Server Configuration (sconfig.exe)**. You will handle many of your initial configuration tasks by using Server Configuration. The screen capture below shows the menu system.

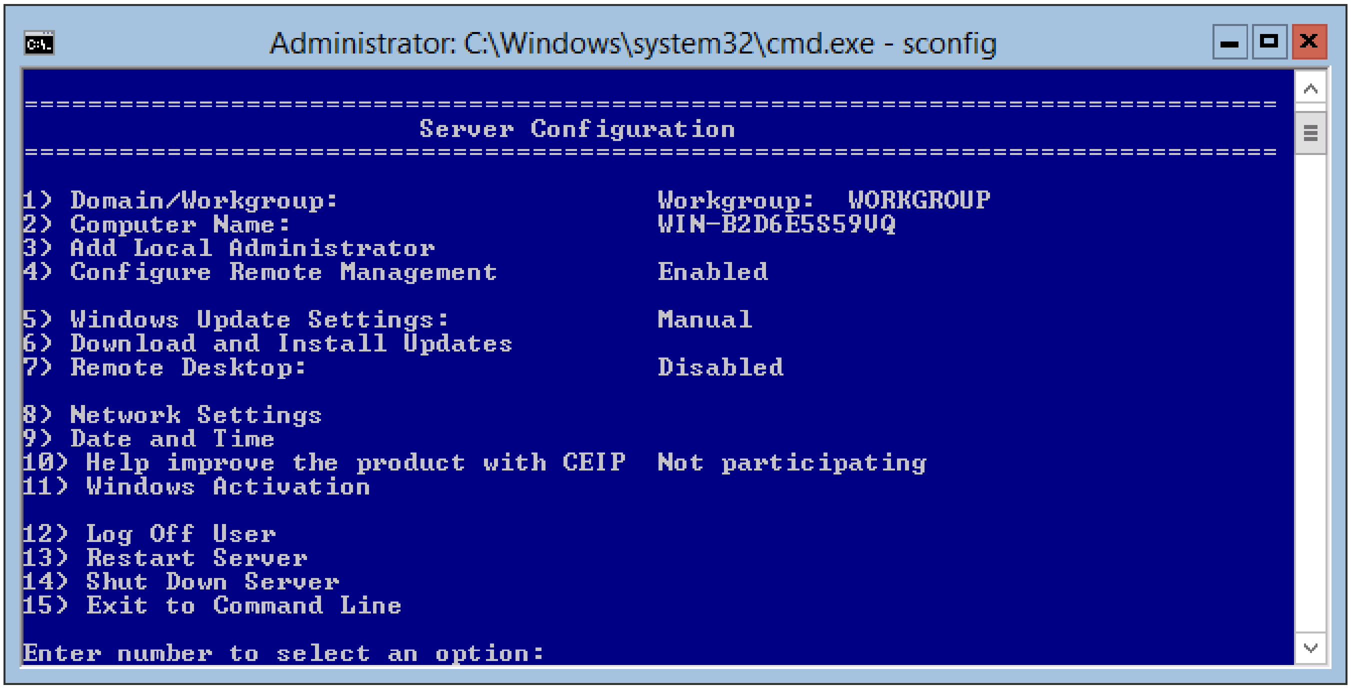

```
===============================================================================
                             Server Configuration
===============================================================================

1) Domain/Workgroup:                    Workgroup:  WORKGROUP
2) Computer Name:                       WIN-B2D6E5S59VQ
3) Add Local Administrator
4) Configure Remote Management          Enabled

5) Windows Update Settings:             Manual
6) Download and Install Updates
7) Remote Desktop:                      Disabled

8) Network Settings
9) Date and Time
10) Help improve the product with CEIP  Not participating
11) Windows Activation

12) Log Off User
13) Restart Server
14) Shut Down Server
15) Exit to Command Line

Enter number to select an option:
```

- **Know how to rename a computer by using PowerShell**. Run the **Rename-Computer** command to rename a computer.
- **Know how to join a domain from PowerShell**. Run the **Add-Computer** command and you will be prompted to enter the domain name and credentials.
- **Know how to enable PowerShell Remoting**. Run the **Enable-PSRemoting** command to enable it.
- **Know how to set up automatic updates**. At an elevated command prompt, navigate to %systemroot%\System32 and then run the following commands in order:
  - **Net stop wuauserv**
  - **Cscript scregedit.wsf /AU 4**
  - **Net start wuauserv**
- **Activate a server**. Run the **slmgr.vbs -ato** command.
- **Work with the firewall**. For example, to disable the firewall on all of the profiles, run the **netsh advfirewall set allprofiles state off** command.

Microsoft put together a web page titled Quick Reference for Server Core Tasks which lists some of the common tasks that you have to perform. Don't memorize all of them. At a minimum, know the tasks in the bulleted list above. If possible, build a Server Core VM in a lab environment and perform some of the configuration tasks.

Delegate administration

For server delegation, there are two local groups that will handle the majority of your delegation needs:

- **Administrators**. When an administrator needs full control of a server, you should add that person to the Administrators group. For anything less, you should look at other options.
- **Remote Desktop Users**. When an administrator needs to connect to a server by using RDP and then run one of the installed applications, you should add that person to the Remote Desktop Users group. In exam items that call out a requirement of minimizing the amount of rights given to administrators, the Remote Desktop Users group may be viable.

You should know how to add users to the groups. Not just by using Computer Management, but also by using the command line. To add a user named Jack in the contoso.com domain to the local Administrators group from the command prompt, run the **net localgroup Administrators /add contoso\jack** command.

Add and remove features in offline images

You use the Dism command to service an offline image. Adding and removing features in offline images is a great topic for exam item writers to use for "build list" type questions where you have to arrange a few or more answers in the correct order.

To add the Telnet Client feature to an offline image located at E:\images\install5.wim, perform the following tasks in order:

- Run the **Dism /Get-ImageInfo /ImageFile:E:\images\install5.wim** command to obtain the name from the output.
- Mount image by running the **Dism /Mount-Image /ImageFile:E:\images\install5.wim /Name:"Corp Image 5" /MountDir F:\temp**.
- Add the feature by running the **Dism /Image:F:\temp /Enable-Feature /FeatureName:TelnetClient** command.
- Commit the changes and unmount the image by running the **Dism /Unmount-Image /MountDir:F:\temp /Commit** command.

Deploy roles on remote servers

You can use Server Manager to add roles on remote servers. But you first need to add the remote servers to Server Manager. Thereafter, when you use the wizard to add a role, you can select from any server that you have added to Server Manager. This is pretty straight forward so I don't expect any exam items to test you on it. However, you should expect to be tested on adding roles on remote servers by using PowerShell.

To add roles or features to a remote server, you need to use the **Install-WindowsFeature -Name <name of role or feature> -ComputerName <hostname of remote server>** syntax. For example, to add Windows Server Backup to Server22, you run the **Install-WindowsFeature -Name Windows-Server-Backup -ComputerName Server22** command. Don't spend any time memorizing the role or feature names because you won't be tested on that knowledge.

Convert Server Core to/from full GUI

From Server Core, you can convert to a full GUI installation by running the **Install-WindowsFeature Server-Gui-Mgmt-Infra, Server-Gui-Shell -Restart** command.

From a full GUI installation, you can convert to the Server Core installation by running the **Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -Restart** command.

Note that in Windows Server 2008 R2 and Windows Server 2008, you had to reinstall the operating system to switch between Server Core installation and the full GUI installation.

Configure services

The primary tasks you perform when configuring services are viewing the status of a service, starting a service, stopping a service, and changing the configuration of a service such as the service credentials. For the exam, you need to know how to use PowerShell to perform the tasks as well as understand some of the advanced configuration settings.

- **View the status of a service with PowerShell**. For example, to view the status of the Remote Desktop Services service, run the **Get-Service TermService** command. To get the names of the services to use in commands, run the **Get-Service** command which outputs all of the services, their name, description, and status.
- **Start a service with PowerShell**. For example, to start the Windows Time service, run the **Start-Service W32Time** command.
- **Stop a service with PowerShell**. For example, to stop the Windows Time service, run the **Stop-Service W32Time** command.

There are other cmdlets related to services too. One in particular, Set-Service, is an excellent distractor that item writers might use because it seems like it would handle any service related configuration. But it doesn't. It can't change the credentials that the service account uses. It also can't change the service recovery which is what action, if any, is taken when the service fails. Instead, you can use sc.exe which is a command line tool that can configuration all aspects of a Windows service. Here are some key commands to know:

- **Change the credentials used for a service**. For example, to change the Windows Time service to use the contoso\w32-svc account with a password of Rdy834h837, run the **sc config obj="contoso\w32-svc" password=" Rdy834h837"** command.
- **Change the recovery of a service**. For example, to set the Windows Time service to restart after a failure, run the **sc failure w32time reset=30 actions=restart/5000** command.

Configure NIC teaming

When you have more than one NIC, you can team them together. Beginning with Windows Server 2012, Windows provides a built-in teaming solution. You can use PowerShell or Server Manager to configure and manage NIC teaming.

The screen capture below shows NIC teaming being configured with Server Manager. To begin, you click Local Server in the left pane of Server Manager and then click the status of NIC Teaming in the right pane (by default, it is "Disabled". Once clicked, the NIC Teaming configuration will be displayed as shown below.

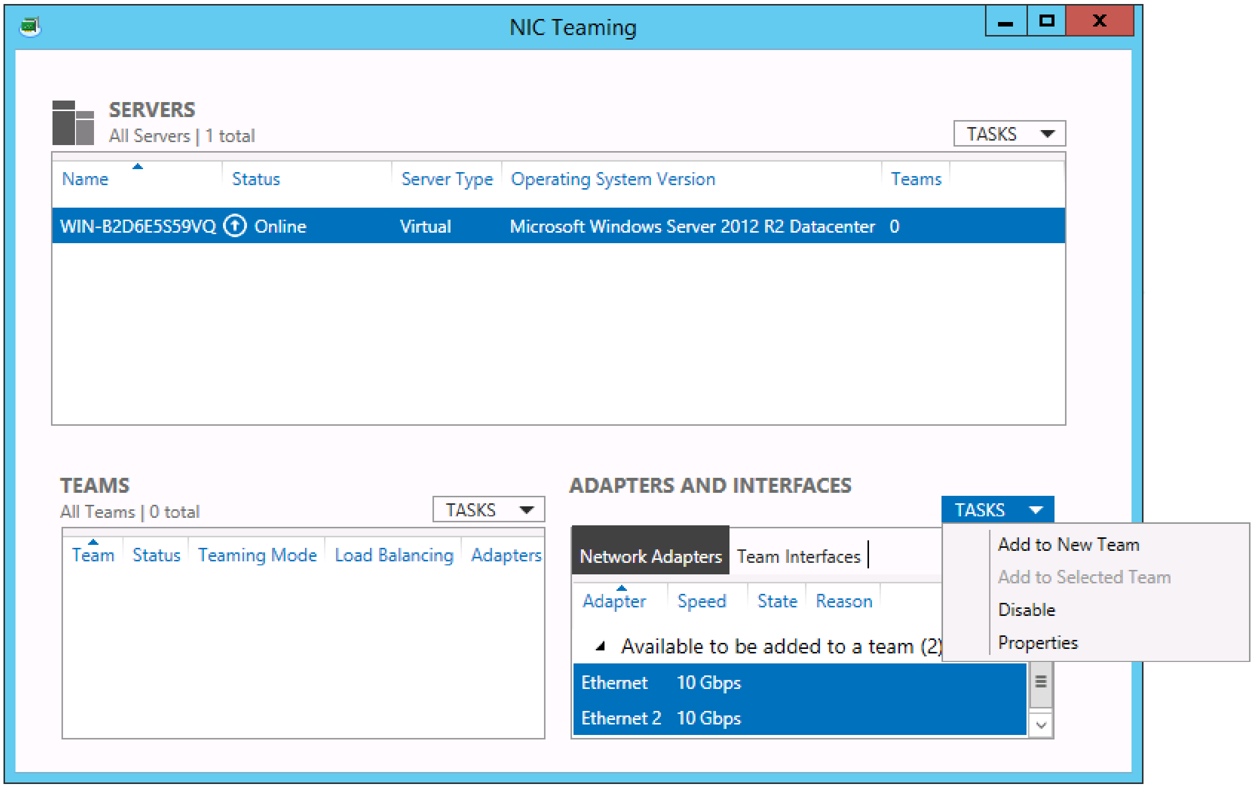

Name	Status	Server Type	Operating System Version	Teams
WIN-B2D6E5S59VQ	Online	Virtual	Microsoft Windows Server 2012 R2 Datacenter	0

Team	Status	Teaming Mode	Load Balancing	Adapters

Available to be added to a team (2)

Adapter	Speed	State	Reason
Ethernet	10 Gbps		
Ethernet 2	10 Gbps		

To create a team, select the NICs, click Tasks, and then click Add to New Team. Then, you name the team. You can configure additional properties at the time of naming, such as whether one adapter is in standby mode.

To create a new NIC team named Team1 with an adapter named “Ethernet” and an adapter named “Ethernet2” by using PowerShell, run the **New-NetLbfoTeam Team1 Ethernet,Ethernet2 -Confirm:$False** command. If you don’t use the -Confirm parameter, the command will ask for confirmation before proceeding.

You should also be aware of the following information for the exam:

- **You can use switch independent mode to handle teaming at the server level**. This is the default setting. In this mode, each NIC is usually connected to a different switch and the switches are not aware of the teaming. Thus, you do not have to configure switches in this mode.
- **You can use switch dependent mode when all NICs are connected to the same switch**. In this mode, you can choose the static method or LACP. LACP offers dynamic link aggregation so use it when supported. In this mode, the switch must be configured too.
- **Use switch independent mode for failover**. If you have two NICs and just want one to act as a standby in case of a failover, use switch independent mode.

Install and configure Windows PowerShell Desired State Configuration (DSC)

DSC is a fairly complicated technology. For administrators that are not familiar with it, getting familiar with it for the exam isn’t a great return on investment. You may need to spend a few hours with it in order to have enough knowledge to answer just a few questions or less. At a high level, DSC is a technology that you can use to consistently install servers and maintain a specified configuration. You can control services, users, groups, applications, registry settings, and more. There are two configuration modes that you can use for DSC:

- **Pull mode**. Pull mode uses a pull server with a web site in IIS or SMB to transmit DSC configuration files to clients. Clients are configured to use the pull server for DSC configuration information and have a scheduled task to perform updates. Pull mode is often the choice for enterprise environments because it simplifies the management process.
- **Push mode**. Push mode is the default mode and it requires that you transmit configuration files to clients on an as needed basis. You can automate the transmission of configuration files. Thereafter, you can run the Start-DscConfiguration command which will immediately indicate the results of the configuration application.

To use DSC, you perform the following high level tasks:

- **Create a Management Object Format (MOF) file**. This file contains the desired configuration. You can use PowerShell to create the file. See this for an example of a .MOF file.
- **Stage the MOF file (Pull mode only)**. This task puts the data on a server named a Pull server which is an IIS server with an OData interface or a server using SMB. In Push mode, there isn’t any staging of the MOF file. Instead, the MOF file is sent upon application (see next step).

- **Enable the MOF file (apply the DSC configuration)**. At this stage, the desired configuration data is pulled (most used) or pushed to the target server(s) for implementation. In Pull mode, a scheduled task applies the DSC configuration, In push mode, you use the Start-DscConfiguration cmdlet to perform this task which will also send the MOF file to the client.

Configure local storage

This section focuses on local storage. When the exam OD was updated for Windows Server 2012 R2, one new topic was added titled “Create storage pools by using disk enclosures”.

Design storage spaces

When you study for the exam, always pay close attention to the verbs used in the exam OD. When you see “Design”, it is often indicative of pre-implementation questions such as when to use it, meeting prerequisites, or figuring out the best way to use it before you implement it. First, know that storage spaces are virtual disks and a configuration. Storage pools, on the other hand, are the physical disks that provide the actual storage.

You can use 3 resiliency types with storage spaces:

- **Simple**. If you use this type, you do not provide any protection against disk failures. If a disk fails, you lose access to the virtual disk. However, if your project requires the highest performance and capacity, this type of storage space may be appropriate. This type of storage is usually used for temporary space. It can have a minimum of one physical disk.
- **Mirror**. This is equivalent to a RAID 1 where at least two disks are writing the same data. It provides redundancy from a single disk failure if you use two disks. If you use five disks, you provide redundancy from two disk failures. However, redundancy comes at a cost. If you have a two-disk mirror made up of one TB disks, you use two TB of space to provide one TB of usable space.
- **Parity**. This is equivalent to a RAID five where data is written across disks in a stripe. It provides redundancy in case of a disk failure. A minimum of three disks is required. To protect against two disk failures, you need seven disks.

Configure basic and dynamic disks

Note that dynamic disks were deprecated in Windows Server 2012. However, one use case – using dynamic disks to create a mirrored volume for the operating system, is still valid and supported. Watch for build list items for configuring disks. This section is well suited for those type of questions because of the strict order that management tasks occur in. For example, let’s review the process of creating a new disk:

1. Add a disk to the computer. Or, if you are using a VM, add a new virtual disk to the VM.
2. Bring the disk online.

3. Initialize the disk. During this process, choose whether to use the MBR or GPT partition type.
4. (Optional step) - Convert to dynamic disk. This is only if you plan to use dynamic disks for software redundancy.
5. Create a volume. During this process, you choose the drive letter, the file system, and some of the other volume configuration options.

For the exam, now the differences between a basic disk and a dynamic disk:

- **Basic disks are the default disks in Windows**. They are the simplest form of storage but don't provide more than generic storage.
- **Dynamic disks offer more features than basic disks**. You can create volumes that span multiple disks (thus creating spanned or striped volumes). You can also use software fault tolerance by opting for mirroring or a RAID 5 volume.

Configure MBR and GPT disks

The MBR partition style is the original partition style available in Windows. The GPT partition style became available with Windows Server 2003 SP1. For the exam, you should know the following information:

- **MBR partition styles support up to 4 partitions**. You can have four primary partitions or three primary partition and one extended partition.
- **MBR partition styles use a single partition table that stores the partition information about the disk**.
- **GPT partition styles support up to 128 partitions**. You need to use a basic disk to create 128 partitions.
- **GPT partition styles support partitions larger than 2 TB**. This is an important fact to remember, especially as disk sizes are growing and large volumes are becoming more common.

Manage volumes

For the exam, you need to know how to create, format, shrink, extend, and delete volumes. These tasks are all straight forward when you use the GUI tools. You should be familiar with the PowerShell and command-line methods as well:

- **Create a new partition**. For example, to create a new partition on Disk 1, run the **New-Partition -DiskNumber 1 -AssignDriveLetter -UseMaximumSize** command.
- **Format a new partition**. For example, to format the E:\ drive with NTFS and label the drive "Data", run the **Format-Volume -DriveLetter E -FileSystem NTFS -NewFileSystemLabel "Data" -Confirm:$False** command.
- **Shrink a partition**. You can do this with the Disk Management console, or you can use the Resize-Partition cmdlet to shrink a partition. When you shrink a volume, Windows scans the volume to

ascertain how much a volume can be safely shrunk. Often, unmovable files limit how much you can shrink a volume.

- **Extend a partition**. You can use Diskpart.exe to extend a partition. However, since Windows Server 2008 R2, you can extend a partition within Windows by using the Disk Management console. Thus, Diskpart isn't needed much with newer versions of Windows Server.
- **Delete a partition**. You can use Diskpart.exe to delete a partition. Or, you can use the Disk Management console. Or, you can use the Remove-Partition cmdlet in PowerShell.

Create and mount virtual hard disks (VHDs)

VHDs are most often used with a virtualization product such as Hyper-V. Thus, when preparing for the exam, it is a good idea to focus on VHDs and their use within Hyper-V. When we talk about a VHD, we are referring to any virtual hard disk (whether using the VHD standard or the newer VHDX standard). There are two primary tasks identified for the exam - creating and mounting.

Before we look at the tasks, there are a couple of important points to know about VHDs for the exam:

- **Do not store a VHD in a folder that is encrypted with EFS**. Hyper-V does not support that. Instead, use BitLocker to protect VHDs.
- **Do not store a VHD in a folder that is compressed with NTFS compression**. It is not supported.

You can create a new VHD by using Hyper-V Manager. This method is straight forward because it uses a wizard-driven process. You can also create a new VHD in Disk Management. You can use PowerShell to create a new VHD too. For example, to create a new 100 GB VHD using the VHDX format, you can run the **New-VHD -Path E:\VMs\Server24\VHDs\Data2.vhdx -SizeBytes 100GB** command. The new .vhdx file will be stored in the specified path. When using PowerShell to create new VHDs, be aware of the following facts:

- **If you do not specify a dynamic, fixed, or differencing disk, then the command will create a dynamic disk**. A dynamic disk does not take up all of the specified space until it is actually used.
- **If you specify a fixed disk during VHD creation, then all of the space specified will be immediately consumed**. A fixed disk takes longer to create but provides slightly higher performance than a dynamically expanding disk.
- **You can specify a differencing disk during VHD creation to use with a VM**. A differencing disk is used to allow seamless roll back to a previous VM state.

The other task for this exam section is focused on mounting VHDs. There are several ways to mount a VHD including the following methods:

- **Use the Disk Management console**. If you right-click Disk Management in the Computer Management console, you have an option to attach a VHD. That is the same as mounting a VHD.

- **Use the Diskpart tool**. The Diskpart tool is an extremely powerful command-line tool. First, you select the disk, then you attach it, then you assign a drive letter.
- **Use PowerShell's Mount-VHD cmdlet**. For example, to mount data.vhdx in E:\data as a read-only disk, run the **Mount-VHD -Path E:\data\data.vhdx -ReadOnly** command. If you omit the -ReadOnly parameter, then it mounts as a readable and writable drive.

Configure storage pools and disk pools

For this section of the exam, start by understanding the order of things. If you want a storage spaces data drive in Windows, the following steps must be performed in order:

1. **Create a storage pool**. A storage pool is a collection of physical disks. The normal term is "storage pool" but "disk pool" can sometimes be used to describe a storage pool. You create a storage pool from disks but you do not bring the disks online or initialize them as part of the process.
2. **Create a virtual disk from the pool**. The virtual disk is the storage space. This wording can be confusing to administrators that are new to storage spaces. You should walk through the creation and configuration in a lab environment as part of your exam preparation because it will help you remember the terminology and order.
3. **Create a new volume from the virtual disk**. You assign a drive letter and can begin using it.

The screen capture below shows a storage pool named Pool1, a virtual disk named VD1 (which is the storage spaces aspect), and the physical disks that make up the backend storage of the storage pool.

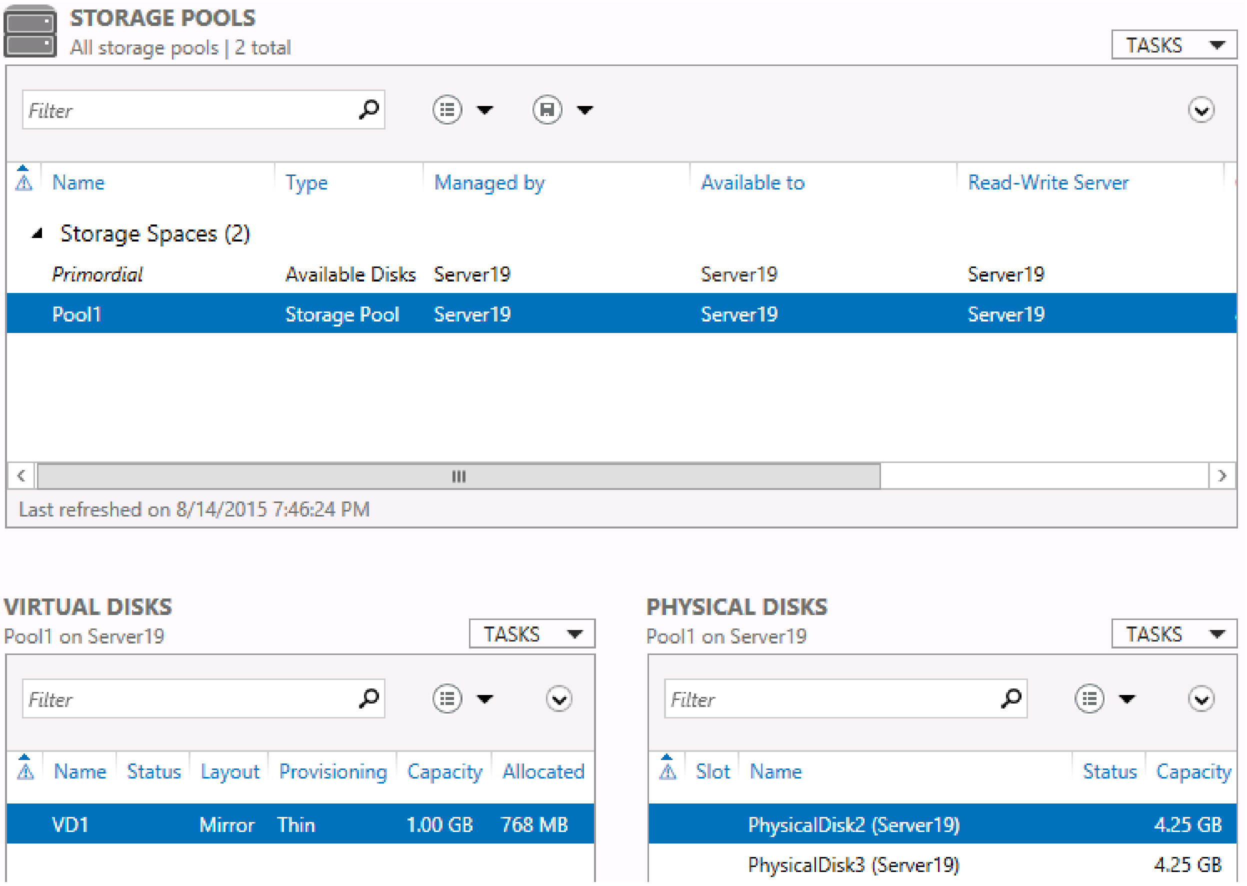

STORAGE POOLS
All storage pools | 2 total

Name	Type	Managed by	Available to	Read-Write Server
Storage Spaces (2)				
Primordial	Available Disks	Server19	Server19	Server19
Pool1	Storage Pool	Server19	Server19	Server19

Last refreshed on 8/14/2015 7:46:24 PM

VIRTUAL DISKS
Pool1 on Server19

Name	Status	Layout	Provisioning	Capacity	Allocated
VD1		Mirror	Thin	1.00 GB	768 MB

PHYSICAL DISKS
Pool1 on Server19

Slot	Name	Status	Capacity
	PhysicalDisk2 (Server19)		4.25 GB
	PhysicalDisk3 (Server19)		4.25 GB

You can create and manage storage pools by using the File and Storage Services console or by using PowerShell. To use PowerShell, you first need to identify all of the physical disks that can be pooled (assuming that you've just added some new disks that you want to pool to the server). To do that, run the **$Disks = (Get-PhysicalDisk -Canpool $True)** command. At that point, the disks are stored in the $Disks variable. Next, run the **New-StoragePool -FriendlyName Pool1 -StorageSubsystemFriendlyName "Storage Spaces*" -PhysicalDisks $Disks** command. PowerShell also offers the Get-StoragePool, Remove-StoragePool, Set-StoragePool, and Update-StoragePool commands.

Create storage pools by using disk enclosures

New to Server 2012 R2 is the support of disk enclosures for storage pools. The important thing to know is that the disk enclosure must not use any type of RAID or fault tolerance and must not abstract the disks in a way that disallows direct access to the physical disks from Windows. If you want to identify disks by slot and take advantage of the disk lights, the enclosure must support SCSI Enclosure Services (SES) version 3.

You can take multiple disk enclosures and use fault tolerance between them. For example, you can create a mirrored storage space with two disk enclosures. Creating a storage pool from disk enclosures is very similar to creating a storage pool from internal physical disks. The primary difference is the new -EnclosureAwareDefault parameter which accepts $True (if you plan to use disk enclosures) or $False (if you don't plan to use disk enclosures).

2 - Configure server roles and features (17%)

This functional group contains a wide range of server technologies include file services and sharing, print server management, and remote management. Some specifics of this functional group include shares and permissions, print drivers and spooling, WinRM, and Windows Firewall. Windows Server 2012 R2 also introduced Work Folders which is mentioned in the exam OD.

Configure file and share access

This objective is primarily focused on file access and control. This includes creating and configuring a file share, configuring how people access those files, working with NTFS permissions and quotas, and configuring Work Folders.

Create and configure shares

Shares can either be single folders or entire volumes that can be accessed by users on a network. You should be familiar with how to create a share by using the GUI so this Study Guide will not cover that. However, you also need to know how to create and configure shares by using PowerShell.

There are a couple of share settings that are not as widely understood as the common share settings and you should spend time to understand them:

- **Simultaneous users**. This setting dictates how many users can access the shared folder at the same time. The default is 16777. For the exam, the most logical way to test this setting is to call out a requirement to reduce the number of simultaneous users (possibly in a scenario where you need to increase performance). Or, in a troubleshooting scenario, a situation where a certain number of users can access a shared folder but whenever another user tries, they cannot gain access. This could be due to the simultaneous users setting being set too low.
- **Caching options**. Caching dictates when and if users can access shared content while offline. You can allow users to dictate which files are available offline. Or, you can specifically configure the shared folder to not allow offline access. The third option is to configure the shared folder so that any files that are opened in the shared folder are automatically made available offline. If you have BranchCache

in your environment, you can optionally enable that with the option to allow users to dictate what is available offline.

To configure a file share by using Windows PowerShell, use the New-SmbShare cmdlet. For example, to configure a file named Share1 for a folder that is located at D:\temp\Folder1, run the following command:

New-SmbShare -Path D:\temp\Folder1 -Name Share1

By default, share permissions are configure with the Read permission granted to the Everyone group. The NTFS permissions of the folder are not modified when you create a share. Watch for scenarios on the exam where you just created a share, assigned Full Control for NTFS permissions, but users are not able to write data to the new share. This can be fixed by adjusting the share permissions.

Configure share permissions

Share permissions are used to define access to a resource over a network. Share permissions do not apply to users that log on locally or log on by using Remote Desktop Services. Share permissions are defined separately from NTFS permissions and are less granular. There are only three levels of share permissions:

- Full Control
- Change
- Read

By default, when you create a share, the Everyone group is granted the Read permission. Additional local or AD DS security groups can be defined to control share access. A good practice for defining share permissions is to define the security groups that will be accessing the share and grant only those groups Change permissions while administrative groups be granted Full Control permissions. By removing the Everyone group from the share permissions, you prevent unauthorized users from even being able to read the data in the share.

Share permissions can also be configured by using the PowerShell cmdlet Grant-SmbShareAccess. The cmdlet configures the same permissions levels but uses different names. The -AccessRight parameter, which is used to grant permissions, accepts the following input:

- Full. This is equivalent to Full Control in the GUI.
- Modify. This is equivalent to Change in the GUI.
- Read. This is the same as in the GUI.

For example, to grant the Marketing group Full Control share permissions on the previously shared directory named Share1, use the following command.

Grant-SmbShareAccess -Name Share1 -AccountName "Contoso\Marketing" -AccessRight Full

Additionally, to remove the permissions that have been granted to an object, use the Revoke-SmbShareAccess cmdlet. For example, to move the Marketing group's access to Share1, run the following command:

Revoke-SmbShareAccess -Name Share1 -AccountName "Contoso\Marketing"

This will remove the object from the access control list on the share, removing all share permissions for the group. Note that revoking the share permissions does not modify any NTFS permissions that might also be defined. Therefore, group members would still be able to access the files locally or through Remote Desktop Services, but not over the network.

Configure offline files

Offline files can be configured to ensure that users can still access the resources stored in a folder without being connected to the network. The most common scenario that this is used is for mobile devices that will still need access to files when they are not connected to a corporate network.

The available caching methods are displayed in the following graphic:

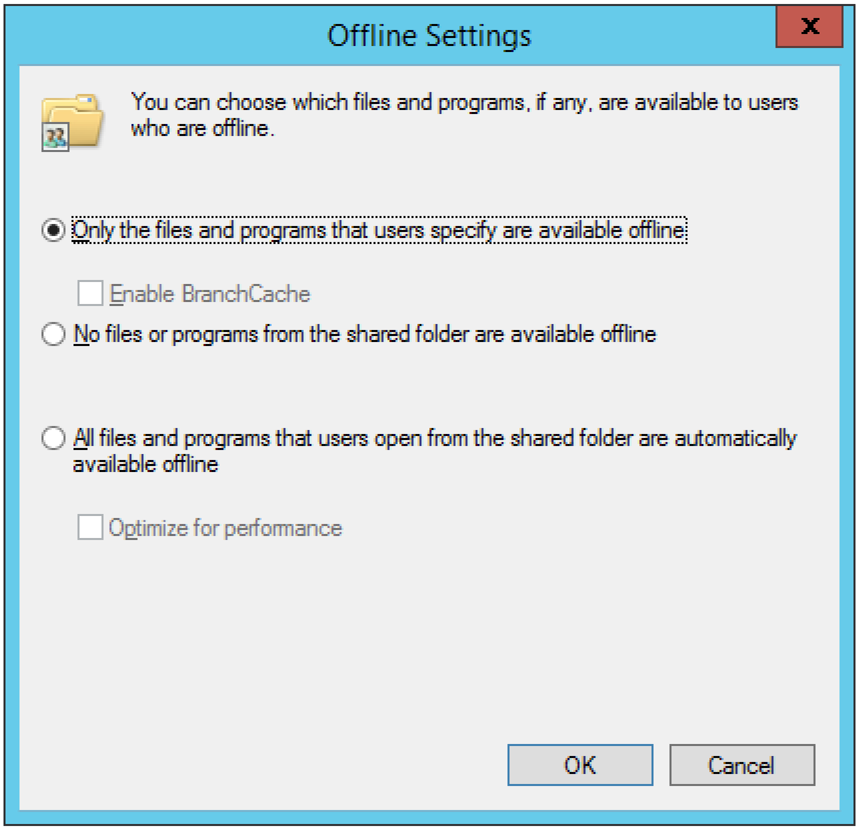

Offline Settings

You can choose which files and programs, if any, are available to users who are offline.

● Only the files and programs that users specify are available offline

□ Enable BranchCache

○ No files or programs from the shared folder are available offline

○ All files and programs that users open from the shared folder are automatically available offline

□ Optimize for performance

- Only files and programs that are specified by users
- No files or programs
- All files and programs

By default, a shared folder is set to allow users to define which files or programs they need offline access to. Additionally, if BranchCache has been configured in the organization, BranchCache can be enabled to enhance the capabilities of offline access.

Configure NTFS permissions

When NTFS permissions are used with Share permissions, the effective permissions are always the most restrictive of the two.

For example, imagine that a folder is shared, and the share permissions are set to grant User1 Full Control share permissions. However, the NTFS permissions for User1 only grant Read permissions. The NTFS permissions are more restrictive, and thus User1 would only be able to read data in the folder. User1 could not write data or modify any files in the folder. Conversely, imagine User2 has been granted Full Control NTFS permissions to the directory but only the share permission of Read. User2 could only read the resources in the directory when accessing the resources over the network. Remember, share permissions are only effective over the network. Because User2 has Full Control NTFS permissions, the permissions would be granted if the user logged on locally or by using Remote Desktop Services.

Configure access-based enumeration (ABE)

ABE can be configured on a shared folder to display or hide shared folders to users based on their NTFS permissions. Note that ABE does not correlate to share permissions. By default, ABE is disabled for folders that are shared. ABE can be enabled by using the Storage and Share Management console, which is part of the File and Storage Services console in Server Manager. Alternately, the New-SmbShare and Set-SmbShare cmdlets can be used with the -FolderEnumerationMode parameter to enable the enumeration mode on a share. The parameter accepts AccessBased to enable enumeration (ABE), or Unrestricted to disable enumeration. For example, if a folder has already been shared with the name Share1, then run the following command to enable enumeration:

```
Set-SmbShare -Name Share1 -FolderEnumerationMode AccessBased
```

For the exam, watch for the following scenarios:

- **An administrator created a new shared folder but users can't see the folder**. In this case, make sure that ABE isn't enabled. Or, check the NTFS permissions that users have to the shared folder.
- **An administrator enabled ABE but users can still see all shared folders**. Check the NTFS permissions and ensure that users don't have indirect access (access via group membership or a group such as Everyone or Domain Users). Ensure that ABE is enabled.

Configure Volume Shadow Copy Service (VSS)

VSS is comprised of three components to perform backups on a system:

- **Requester**. The requestor is the source of the request for a shadow copy.
- **Provider**. The provider is the interface to create and maintain a shadow copy.
- **Writer**. The writer ensures that data is ready before creating a shadow copy.

The requestor initiates the backup process or restore process. The provider manages the process by instructing the writer which tasks to perform.

VSS was updated in Windows Server 2012 to also manage tasks over a remote SMB 3.0 share. This capability extends to UNC paths, including shares on external NAS devices. To use this capability, the environment must meet the following requirements:

- The application and file servers being backed up must be running Windows Server 2012 or later and be joined to the same AD DS domain.
- The File Server VSS Agent Service role service must be installed on the application and file server that is being backed up.
- The File Share Shadow Copy agent must be a member of the Backup Operators or Administrators groups on both the application server and the file server that is being backed up, and have at least read permissions on the data to be backed up.

For more information on using VSS with SMB shares, see this.

Configure NTFS quotas

NTFS supports disk quotas on local volumes or shared folders. On volumes, NTFS quotas can be enabled to warn or deny users when writing data to the volume. Additionally, alerts can also be generated to the event log. For shared folders, FSRM can be used to apply a quota template to a shared folder. Templates have pre-defined settings for the limit, threshold, and alerting options. You can create customized templates by using the FSRM GUI or by using the **New-FsrmQuotaTemplate** cmdlet. For example, to create a new template that limits a shared folder size to 5 GB, run the following command:

```
New-FsrmQuotaTemplate -Name "5 GB Limit" -Size 5GB
```

You can apply a quota to shared folder without using a template by using the New-FsrmQuota cmdlet. For example, to create a quota of 5 GB to a single folder named Share1 in the D:\temp folder, run the following command:

```
New-FsrmQuota -Path "D:\temp\Share1" -Size 5GB
```

For more information on quotas, see https://technet.microsoft.com/en-us/library/cc938945.aspx.

Create and configure Work Folders

Work Folders is a technology introduced with Windows Server 2012 R2. It enables users to synchronize corporate file shares with any supported devices. You should famliarize yourself with the prerequisites:

- To use Work Folders, the file server must be running Windows Server 2012 R2 and the shared folders must be on an NTFS volume.
- Work Folders are supported with Windows 7 and later. However, Windows 7 computers require a separate download.
- Windows 7 computers must be excluded from the Work Folders password policies (if in use). Instead, you need to use Group Policy to enforce password policies on Windows 7 computers.

- For all file servers, a server certificate must be used and it must be issued from a trusted Certification Authority.
- Windows 7 Professional or higher, Windows 8.1 and Windows RT 8.1, and iOS 8 or later are supported operating systems.

Devices must have enough free space to store the contents of the Work Folder locally, plus an additional 6 GB. By default, Work Folders uses the %USERPROFILE%\Work Folders location. The additional 6 GB is not required if the default location is moved to a different volume. The maximum individual file size that can be synchronized is 10 GB. However, folder quotas can be configured by using FSRM.

For additional information, see https://technet.microsoft.com/en-us/library/dn265974.aspx.

Configure print and document services

This objective covers the primary functions of defining network printing services, from drivers, spool management, priorities and permissions.

Configure the Easy Print print driver

In a Remote Desktop Services environment, the RD Session Host will attempt to use the Easy Print driver when printing to a locally-defined printer. This setting can be disabled by using Group Policy, which will cause the RD Session Host to attempt to locate a suitable driver for the printer. If a driver is not available, then the Easy Print driver will still be used for the printer.

Configure Enterprise Print Management

Print Management requires the Print and Document Services role to be installed on the server that you plan to manage printers from. To perform print management tasks, the managing user account must be a member of one of the following groups:

- Print Operators
- Server Operators
- Local administrators (which includes Domain Admins by default)

In addition to the appropriate permissions, the Windows Firewall must have the Print Management exception added for remote management.

Configure drivers

Print servers must have drivers installed for all clients that will be printing to them. For example, a print server running Windows Server 2012 with 64-bit Windows 8 and 32-bit Windows 7 clients must have both versions of the print drivers installed. If the installation of 32-bit drivers fail on the server, attempt to install the drivers on the print server from the client.

Drivers can be updated by simply installing the latest version of the driver, and then selecting the driver for the appropriate printer. You can remove a driver by selecting delete from the list of available drivers in the GUI.

Configure printer spooling

The print spooler can be used to control the communication between the print server and the printer. By using the spooler, an application can release control of the print job sooner. There are many options that can be configured by using the spooler, including:

- start printing immediately
- start printing after the last page is spooled
- hold mismatched documents
- print spooled documents first
- keep documents after they are printed

Spooling can be configured to print immediately or after the last page is spooled. Printing the document immediately ensures that the print job finishes faster, but could delay other jobs if the print job is large. Holding mismatched documents will ensure that jobs will not print if they do not match the configuration of the printer. For example, a job that specifies an A4 paper type will not print on a printer that only has letter size paper. Keeping documents after they are printed ensures that future print jobs are printed faster but can use large amounts of disk space.

Configure print priorities

Priorities can be configured to determine in which order a printer or job is selected to complete printing. Printers and print jobs can be configured with a priory from 1 to 99, with 1 being the lowest priority, and 99 the highest. Priorities only affect the next print job and do not modify jobs that have already begun printing. For example, a printer has the print jobs displayed in the table below.

Job Number	Status	Job Priority
1	Printing	1
2	Spooled	5
3	Spooled	1
4	Spooled	99
5	Spooled	5

In this example, while the active job has the lowest priority, it will still complete because it has already started. Assuming that no other jobs will be submitted, the jobs will be printed in the following order:

Job Number	Reasoning
1	Actively printing
4	Has the highest priority of spooled jobs
2	Next-highest priority and submitted earlier than job 5
5	Next-highest priority
3	Lowest priority

The priority of a print job can be configured from the properties of the print job.

Configure printer permissions

Printer permissions are managed similar to share and NTFS permissions but have the following options for users and groups:

- Print
- Manage this printer
- Manage documents
- Special permissions

By default, the Everyone security principal is granted Print permissions when a printer is added. Administrators are granted additional permissions to manage the printer and the documents submitted to be printed. The Manage this printer permissions must be granted to perform printer-specific tasks such as configuring spooler options. The Manage documents permission must be granted to perform job-related tasks such as changing job priority. Special permissions include Take Ownership and Change Permissions.

Configure servers for remote management

This objective covers configuring servers for remote management including for day to day tasks. Server Core and the Windows Firewall are topics within this objective. Also included is managing servers that are both domain-joined and non-domain joined.

WinRM

Windows Remote Management, or WinRM, is a remote management component for Windows servers. In Windows Server 2012 and later, WinRM is enabled by default. In earlier versions of Windows, WinRM must be enabled, for example, by running the **winrm quickconfig** command. You can view whether the WinRM service is currently running by using the following PowerShell command:

Get-Service WinRM

Configure down-level server management

Down-level server management is the practice of managing earlier versions of Windows from Windows Server 2012 R2. Windows Server 2008 and 2008 R2 can be managed by using the Server Manager GUI, or by using Windows PowerShell. Remember, to manage a remote server with a specific role, the tools must be installed on the management server. However, Server 2008 and Server 2008 R2 do not have WinRM enabled by default. Run the **winrm quickconfig** command to enable remote management.

Configure servers for day-to-day management tasks

Day to day management should focus on tasks that make remote management easier or more intuitive. This objective may focus on enabling remote disk management and event log management. Familiarize yourself with the Enable-NetFirewallRule cmdlet to manage the Windows Firewall and to enable remote management of specific services.

Configure multi-server management

With WinRM and the Server Manager console, a single server can be configured to manage all servers in an environment. Simply use the Server Manager console to add other servers in an environment to the GUI. You can then manage any server, role, or role service from the Server Manager console, provided the management binaries have been installed on the local server. For example, a management server can have the FSRM management tools installed and manage all of the file servers.

Configure Server Core

A Server Core installation can be managed by traditional management tools, including:

- PowerShell
- Server Manager
- Microsoft Management Console

Alternately, if an application requires a graphical interface to operate correctly, a GUI can be added to a Server Core installation by using the Install-WindowsFeature cmdlet. For example, to install the GUI by using Windows Update, run the following command:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart
```

Configure Windows Firewall

The Windows Firewall can be managed remotely by using the Windows Firewall with Advanced Security MMC snap-in or by using PowerShell. The destination computer must have WinRM enabled and the IPSec and Windows Firewall services must be running.

Manage non-domain joined servers

In addition to WinRM being enabled, non-domain joined servers must also be added to the trusted hosts list of WinRM. For example, to add a computer named Server1 to the list of computers that can manage the non-domain joined server, run the following command on the non-domain joined server:

```
winrm s winrm/config/client '@{TrustedHosts="Server1"}'
```

Non-domain joined servers can easily be managed by using Windows PowerShell. PowerShell provides a built-in method of specifying alternate credentials by using the Get-Credential cmdlet, which would need to be used to authenticate with the non-domain joined server. When specifying the credentials, you may need to specify that they are local. For example, if a workgroup computer named Server1 has an administrative user named User1, the following format should be specified for the user name:

- Server1\User1
- Localhost\user1
- .\User1

This syntax ensures that User1 will be authenticated as a local user on the remote server. The New-PSSession cmdlet can be used to generate a new PowerShell session with a specific server and specify the credentials for the local server. For example, to connect to a workgroup server named Server1, run the following commands:

$creds = Get-Credential

New-PSSession -ComputerName Server1 -Credential $creds

3 - Configure Hyper-V (18%)

This functional group covers configuring a virtual machine for use, including VM settings, storage settings, and network settings.

Create and configure virtual machine settings

This objective covers the basic functionality of creating a virtual machine with Windows Server 2012 R2 and Hyper-V, as well as configuring the basic properties of a virtual machine during and after creation.

Configure dynamic memory

By default, when using dynamic memory with Windows Server 2008 or later, the machine will be granted 512 MB of RAM at startup. The amount of RAM that is assigned at startup, as well as a maximum RAM limit, can be configured by using the Set-VMMemory cmdlet. For example to enable dynamic memory on a VM named VM1 with 512 MB of RAM at a minimum, a startup value of 1 GB, and a maximum limit of 6 GB, run the following command:

Set-VMMemory VM1 -DynamicMemoryEnabled $True -MinimumBytes 512MB -StartupBytes 1GB -MaximumBytes 6GB

Altnernatively, Hyper-V Manager can be used by modifying the properties of the virtual machine.

Configure smart paging

Hyper-V uses smart paging in the following scenarios:

- when a VM is being restarted (not powered on)
- if there is no available physical memory
- if memory cannot be reclaimed from other VMs

You can configure the location of the Smart Paging file by using the Set-VM cmdlet. For example, to set the file path to D:\Temp on VM named VM1, run the following command:

```
Set-VM -Name VM1 -SmartPagingFilePath “D:\Temp”
```

Configure Resource Metering

Resource Metering can be configured to collect metrics on virtual machines in Hyper-V. Resource Metering can be enabled for a specific virtual machine by using the **Enable-VMResourceMetering** cmdlet. For example, to enable Resource Metering on a VM named VM1, run the following command:

```
Enable-VMResourceMetering –VMName VM1
```

Additionally, the metrics can be filtered by including the **-ResourcePoolType** parameter. This enables you to collect metrics on VM components including:

- Processor
- VHD
- Ethernet
- Memory

Configure guest integration services

Most guest integration services are enabled by default, and are available for both Windows Server 2012 and Windows Server 2012 R2. However, the Guest Services functionality is limited to Windows Server 2012 R2. Guest Services is disabled by default. Guest Services functionality enables you to copy files to a VM without using a network connection. To copy a file to a VM by using the integration services, use the **Copy-VMFile** cmdlet. For example, to copy an .iso file named en_sql_server_2014_enterprise_edition_x64_dvd_3932700.iso that is located on the D:\ of the host, to D:\Temp on the VM, run the following command:

```
Copy-VMFile "TestVM" -SourcePath "D:\ en_sql_server_2014_enterprise_edition_x64_dvd_3932700.iso" -
DestinationPath "D:\Temp\ en_sql_server_2014_enterprise_edition_x64_dvd_3932700.iso " -
CreateFullPath -FileSource Host
```

Create and configure Generation 1 and 2 virtual machines

Windows Server 2012 R2 introduced a new generation of virtual machines that can be created. Generation 2 VMs support the following features

- PXE boot by using a standard network adapter
- Boot from SCSI VHD or virtual DVD
- Secure boot (enabled by default)
- UEFI firmware support

The generation selection can be made during virtual machine creation. Generation 1 is the default for new VMs. The following operating systems can be installed on Generation 2 VMs:

- Windows 8 (64-bit)
- Windows 8.1 (64-bit)
- Windows Server 2012
- Windows Server 2012 R2

After a VM has been created, the generation cannot be changed. Some Linux distributions are supported on Generation 2 VMs, but secure boot must be disabled. For more information on Generation 2 VMs, see https://technet.microsoft.com/library/dn282285.aspx.

Configure and use enhanced session mode

Enhanced session mode enables you to use Remote Desktop to connect to virtual machines through the VMbus, which does not require a network connection. Therefore, you can connect to VMs and use the features offered by the Remote Desktop Protocol, while ensuring the security or network segmentation that a VM might need. When you connect to a VM from the Hyper-V Manager, you will be prompted for the screen resolution to use for the Remote Desktop connection. Other resources such as disk drives, clipboard, sound, and printers can be passed through the Remote Desktop connection.

Create and configure virtual machine storage

This objective covers configuring the storage aspects of a virtual machine, including the disk type, location, and method of accessing the available storage.

Create VHDs and VHDX

The two types of virtual disks are VHD and VHDX. VHDs support sizes up to 2 TB, and can be used with any supported operating system. VHDX disks support sizes up to 64 TB, but can only be used with Windows 8 and Windows Server 2012 and later operating systems.

A VHD or VHDX can be in three different formats:

- Fixed virtual hard disk
- Dynamically expanding virtual hard disk
- Differencing virtual hard disk

A new disk can be created by using the wizard in the Hyper-V manager or by using the **New-VHD** cmdlet. When using the wizard, the default disk type is Fixed when VHD is selected, and Differencing when VHDX is selected. When using the cmdlet, you must specify the disk format by using the following parameters:

- -Fixed
- -Dynamic
- -Differencing

For example, to create a fixed VHDX disk in D:\temp named Disk1 with a sized of 500 GB, run the following command:

```
New-VHD -Path D:\temp\Disk1.vhdx -SizeBytes 500GB
```

Configure differencing drives

A differencing disk should not be used in production environments but can be used for labs and test environments. A differencing disk is associated with a parent disk, and all new writes or changes are written to the differencing disk. The parent disk is only used to read content and is not written to unless the differencing disk is merged with the parent disk. Differencing disks act like dynamically expanding disks growing in size with each additional write. Differencing disks do not shrink when data is deleted from the disk.

To create a differencing disk at D:\temp\Diff1.vhdx that has a parent disk of D:\temp\Parent.vhdx, run the following command:

```
New-VHD -ParentPath D:\temp\Parent.vhdx -Path D:\temp\Diff1.vhdx -Differencing
```

Modify VHDs

There are several options to modify a VHD or VHDX, including:

- Compact
- Convert
- Expand

Compacting a VHD or VHDX enables you to regain space on the host without modifying the storage capacity of a disk. Typically, this is only necessary if you delete a large amount of files from the VHD and need to regain the space on the host system. This operation can also be performed by using the **Optimize-VHD** cmdlet. For example, to compact a disk located at D:\temp\VM1.vhdx, run the following command:

Optimize-VHD -Path D:\temp\VM1.vhdx -Mode Full

Acceptable modes include:

- -Full
- -Pretrimmed
- -Prezeroed
- -Quick
- -Retrim

For VHDs, the default option is Full. For VHDXs, the default option is Quick.

Converting a VHD or VHDX presents you with both the options of converting to disk type, such as VHD to VHDX or VHDX to VHD, as well as modifying the disk format. A dynamically expanding disk can be converted into a fixed disk, or a fixed disk can be converted into a dynamically expanding disk. A differencing disk cannot be modified. When converting a disk, the underlying operation copies the data and creates a new disk with the selected properties. Converting a VHD can also be performed by using the Convert-VHD cmdlet. For example, to convert a VHD to VHDX located at D:\tepmVM1.vhd, run the following command:

Convert-VHD -Path D:\temp\VM1.vhd -DestinationPath D:\temp\VM1.vhdx

The -Expand option enables you to expand the VHD or VHDX up to the maximum supported size for the disk format. For VHD disks, this is 2 TB, and for VHDX, it is 64 TB. Expanding a virtual disk can also be performed by using the Resize-VHD cmdlet. For example, to expand D:\temp\VM1.vhdx to 1 TB, run the following command:

Resize-VHD -Path D:\temp\VM1.vhdx -SizeBytes 1TB

Configure pass-through disks

Storage disks, including storage presented from a SAN, can be presented to a VM through Hyper-V. Pass-through disks can either be presented as-is or converted to be a VHDX presented to a VM. Either configuration can be configured through the Hyper-V Manager. Alternatively, the Add-VMHardDiskDrive cmdlet can be used to add physical disks. For example, to add Physical Disk 2 to a VM named VM1, on SCSI Controller 0, run the following command:

```
Add-VMHardDiskDrive -VMName VM1 -ControllerType SCSI -ControllerNumber 0 -DiskNumber 2
```

Manage checkpoints

Checkpoints (sometimes called snapshots), enable you to create a point in time reference of a VM's disks. A checkpoint can be created through Hyper-V Manager or by using the Checkpoint-VM cmdlet. For example, to create a checkpoint of a VM named VM1 and name the checkpoint BeforeApplicationInstall, run the following command:

```
Checkpoint-VM -Name VM1 -SnapshotName BeforeApplicationInstall
```

Checkpoints should never be used as a replacement for backups because they do not use a VSS provider. Hyper-V limits a virtual machine to a maximum of 50 checkpoints.

Implement a virtual Fibre Channel adapter

A Fibre Channel Adapter with editable port addresses can be added to a virtual machine by using Hyper-V Manager or by using the Add-VMFibreChannelHba cmdlet. When adding a Fibre Channel HBA to a VM, it must be associated with a Virtual FC SAN similar to how a NIC is associated with a switch. The Virtual SAN is associated with physical HBAs on the Hyper-V host. To add a FC HBA to a VM named VM1, and a virtual SAN named SAN1, run the following command:

```
Add-VMFibreChannelHba -VMName VM1 -SanName SAN1
```

Configure storage Quality of Service (QoS)

Quality of Service for storage devices is a new features with Windows Server 2012 R2. Setting a QoS policy for a virtual machine can be accomplished by using the Settings menu of a VM, as displayed in the following screen capture.

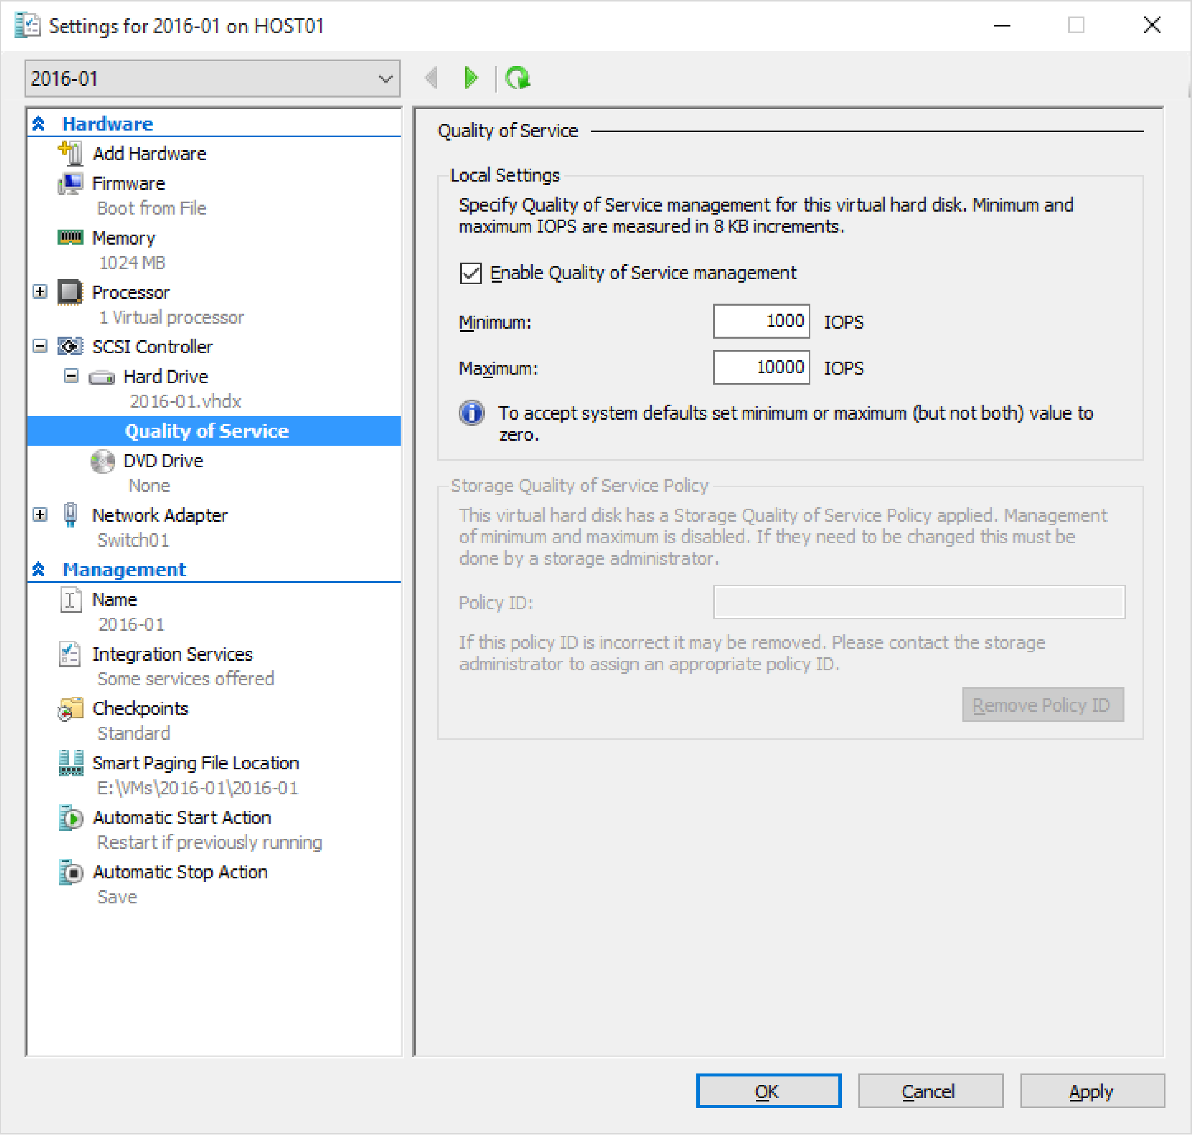
Settings for 2016-01 on HOST01

2016-01

Quality of Service

Local Settings

Specify Quality of Service management for this virtual hard disk. Minimum and maximum IOPS are measured in 8 KB increments.

☑ Enable Quality of Service management

Minimum: 1000 IOPS

Maximum: 10000 IOPS

To accept system defaults set minimum or maximum (but not both) value to zero.

Storage Quality of Service Policy

This virtual hard disk has a Storage Quality of Service Policy applied. Management of minimum and maximum is disabled. If they need to be changed this must be done by a storage administrator.

Policy ID:

If this policy ID is incorrect it may be removed. Please contact the storage administrator to assign an appropriate policy ID.

Alternatively, the minimum and maximum IOPS for a disk can be set by using the Set-VMHardDiskDrive cmdlet. For example, to set the maximum IOPS of the SCSI disk associated with VM1 to 5000 and the minimum IOPS to 500, run the following command:

```
Set-VMHardDiskDrive -VMName VM1 -ControllerType SCSI -MaximumIOPS 5000 -MinimumIOPS 500
```

Create and configure virtual networks

This objective covers how a VM connects to other VMs and physical networks. It defines virtual switches, performance, and VM network settings.

Configure Hyper-V virtual switches

Virtual switches are the communication method between the VMs and depending on the virtual switch type, the physical network adapters. There are three switch types:

- **External switch**. An external switch binds to the physical network adapter so that the VMs that have a NIC connected to the switch can communicate with other devices on the physical network.
- **Internal switch**. An internal switch provides an internal connection between the VMs that are running on the host and the host itself. An internal switch does not bind to any physical adapters and cannot communicate with other devices on a physical network.
- **Private switch**. A private switch provides an internal connection between the VMs that run on the physical host only. The switch does not provide any communication from the VMs to the host or to the devices on a physical network.

Only external and internal networks can define a VLAN ID when configuring the switch. Private networks cannot have VLANs configured. A virtual switch can be created by using the New-VMSwitch cmdlet. For example, to create a private network named Private-HQ, run the following command:

```
New-VMSwitch "Private-HQ" -SwitchType Private
```

Optimize network performance

Bandwidth management can be configured on a per-VM basis. The network adapter properties in the settings of the VM enables you to set minimum and maximum bandwidth values. Additionally, you can use the Set-VMNetworkAdapter cmdlet with the following parameters:

- -MaximumBandwidth
- -MinimumBandwidthAbsolute
- -MinimumBandwidthWeight
- -AllowTeaming

These parameters enable you to configure a network adapter to optimize the performance of the VM. For example, to set a maximum bandwidth of 100 Mbps for the network adapter on a VM named VM1, run the following command:

Set-VMNetworkAdapter -VMName VM1 -MaximumBandwidth 100000000

Configure MAC addresses

By default, MAC addresses for VM network adapters are set to dynamically assigned. VM NICs can also use defined static MAC addresses or have MAC address spoofing enabled. MAC addresses can be configured by using the Set-VMNetworkAdapter cmdlet. For example, to set a static MAC address of 00:15:3D:01:09:10 on a VM named VM1, run the following command:

Set-VMNetworkAdapter -VMName VM1 -StaticMacAddress 00:15:3D:01:09:10

You can use the -DynamicMacAddress parameter to switch the adapter back to a dynamic address.

Configure network isolation

There are four options for configuring network isolation:

- None
- NativeVirtualSubnet
- ExternalVirtualSubnet
- VLAN

To configure an individual adapter, use the Set-VMNetworkAdapterIsolation cmdlet. For example, to configure the network adapter on a VM named VM1 to use VLAN isolation, run the following command:

Set-VMNetworkAdapterIsolation -VMName VM1 -IsolationMode VLAN

Configure synthetic and legacy virtual network adapters

By default, Generation 1 and Generation 2 VMs use synthetic virtual network adapters. For Generation 1 VMs, a legacy network adapter type can be added after VM creation by using the GUI or PowerShell. A legacy network adapter is typically used to perform a network-based installation of a guest operating system, although this may vary depending on how the VM was deployed.

You can add a legacy network adapter by using the Add-VMNetworkAdapter cmdlet. For example, to add a legacy network adapter to a VM named VM1 on the virtual switch named Switch1, run the following command:

Add-VMNetworkAdapter -VMName VM1 -IsLegacy $True -SwitchName Switch1

Configure NIC teaming in virtual machines

NIC teaming can be enabled for VMs by using multiple virtual network adapters. Teaming must be enabled on the virtual NIC for use by the guest VM. This is typically only configured if the Hyper-V server does not have multipathing or teaming already configured. To configure NIC teaming, use the Set-VMNetworkAdapter cmdlet. For example, to enable NIC teaming on the network adapters for a VM named VM1, run the following command:

```
Set-VMNetworkAdapter -VMName VM1 -AllowTeaming On
```

The acceptable options for the -AllowTeaming parameter are On or Off. Yes/No and $True/$False are not accepted by the parameter.

4 - Deploy and configure core network services (17%)

This functional group covers a wide range of networking technologies including IP addressing, DHCP, and DNS. Some of the topics are well known by many administrators. Others, such as configuring Teredo and configuring ISATAP, are lesser known. For this functional group, there weren't any changes when the objective domain (OD) was updated for Windows Server 2012 R2. This often is indicative of there not being a lot of major changes in Windows Server 2012 R2 for the functional group technologies.

Configure IPv4 and IPv6 addressing

This section is all about IP addressing. Included is subnetting which many administrators loath. Also included is interoperating with IPv4 and IPv6. Between those two, there are many areas to focus on to prepare for the exam.

Configure IP address options

IP address options include the following items:

- IP address
- Subnet mask
- Gateway
- DNS servers
- Suffix search order list

You should be familiar with configuring IP address options by using DHCP, Netsh, and PowerShell. Many administrators are comfortable with DHCP but less so with PowerShell and Netsh.

- **Get-NetIPAddress**. Use this PowerShell command to get the current IP address information on a computer. Often, you will use this to obtain the values of the InterfaceIndex and InterfaceAlias properties. You use these properties in configuration commands.
- **Set-NetIPAddress**. Use this PowerShell command to add an IP address to a computer. You need to specify the desired InterfaceIndex or InterfaceAlias to use this command.
- **Netsh**. To add an IP address to the Ethernet interface, you would run the **netsh interface ip add address "Ethernet" 10.10.10.100 255.255.255.0** command. There are a lot of ways to use NETSH to manage IP address settings. Be familiar with adding and removing addresses and changing a computer from static to DHCP or from DHCP to static. See https://technet.microsoft.com/en-us/library/cc738592(v=ws.10).aspx for additional detail on using NETSH for managing IP address settings.

Configure subnetting

This section on configuring subnetting is likely to be focused on a single topic – how many computers you can fit into a given subnet based on the subnet mask. You should be able to figure out the appropriate subnet mask to use if you need to put a specified number of computers into a single subnet. It is helpful to know the math. At worst, memorize some of the most popular subnet masks and the number of computers that fit into the subnets by studying the table below.

Subnet mask	Prefix	Max number of computers
255.255.255.0	/24	254
255.255.254.0	/23	510
255.255.252.0	/22	1022
255.255.248.0	/21	2046
255.255.240.0	/20	4094
255.255.224.0	/19	8190
255.255.192.0	/18	16382

Configure supernetting

For the exam, you'll need to know only the basics of supernetting. Back in the early days of the internet, there were three primary networks (often called a classful network):

- **Class A network**. A Class A network is the largest network and can have 16,277,214 hosts. In the early days of the internet, large organizations such as Ford Motor Company were assigned a Class A network of public IP addresses.
- **Class B network**. A Class B network can have up to 65,534 hosts. Like Class A networks, many companies were assigned a Class B network of public IP addresses.
- **Class C network**. A Class C network can have up to 254 hosts.

Classless interdomain routing (CIDR) was introduced in the early 1990s to move away from the rigidity of classful networking. It wasn't immediately implemented though. Instead, it wasn't used on a wide scale until the internet began to run low on public IP addresses. As the internet grew in popularity, providers turned to CIDR to reduce waste. CIDR has specific notation which most administrators are already familiar with. Instead of calling a network that can hold up to 254 hosts a Class C network, you would say it is a /24 network. With CIDR, if a company needs 1,000 public IP addresses, they can be assigned a /22 which can have up to 1,022 hosts. In the early days of the internet, that company would often be assigned a Class B network instead.

For the exam, you should understand how to break up a classful network into smaller networks. For example, you might be presented with a scenario where you have a Class C network and you need to break it up into 2 networks. You should know how to compute the subnet mask. The tables below show some of the common networks that can be formed from Class C and Class B networks.

Subnet mask	**Prefix**	**Max # of computers**	**Comment**
255.255.255.128	/25	126	Split Class C in half
255.255.255.192	/26	62	Split Class C in 4
255.255.255.224	/27	30	Split Class C in 8
255.255.255.240	/28	14	Split Class C in 16

Subnet mask	**Prefix**	**Max # of computers**	**Comment**
255.255.128.0	/17	32766	Split Class B in half
255.255.192.0	/18	16382	Split Class B in 4
255.255.224.0	/19	8190	Split Class B in 8
255.255.240.0	/20	4094	Split Class in in 16

Configure interoperability between IPv4 and IPv6

Microsoft introduced IPv6 support in Windows Vista and Windows Server 2008. As part of that introduction, new technologies were implemented to allow computers using IPv6 to function on a network (or on the internet) running IPv4. For the exam, you should familiarize yourself with ISATAP and Teredo as well as 6to4.

Configure ISATAP

ISATP facilitates IPv6 communication between hosts across an IPv4 intranet. For the exam, remember that this technology is for an intranet and not the internet. For ISATAP functionality across the internet, 6to4 is the technology to use. It isn't called out specifically for this exam, but it may come up in some scenarios. For more information, see the IPv6 Transition Technologies whitepaper at http://download.microsoft.com/download/1/2/4/124331bf-7970-4315-ad18-0c3948bdd2c4/IPv6Trans.doc. However, only review bits and pieces for the exam. If you try to consume the entire 47 pages in detail, the return on the investment for the 70-410 exam will be very low.

Configure Teredo

Teredo is an IPv6 transition technology. It facilitates the communication of two computers running IPv6 across IPv4 NATs (and across the internet). A Teredo client is included with all Windows operating systems since Windows XP SP2 and Windows Server 2003 SP1. Historically, Teredo was used more often than 6to4 due to a lack of support for 6to4 in edge devices. But Microsoft's stance is that Teredo is a last resort technology with the eventual goal that it will not be used at all (instead, ISATAP or 6to4 will be used). For more information on Teredo, see this. However, don't spend a ton of time on it. To really learn the intricacies of Teredo can take weeks. For the exam, you should spend a couple of hours for the section on configuring IP addressing.

Deploy and configure Dynamic Host Configuration Protocol (DHCP) service

DHCP has several topics on the exam. Based on the percentage of the exam that the functional group takes up and the number of topics, I expect that most test takers won't see a question on all of these DHCP topics. But I expect that some might. From an exam preparation perspective, install and configure DHCP in a lab environment as part of your preparation. This will allow you to run through some of the basic DHCP topics listed below.

Create and configure scopes

Be familiar with the DHCP management console. This study guide won't spend much time there with the assumption that most administrators are familiar with it already. However, you should also know how to create and configure scopes using PowerShell and the command line, as follows:

- Use PowerShell to create a new scope named Scope1 with a range of 192.168.1.100 to 192.168.1.200 in a /24 network:

  Add-DhcpServerv4Scope -Name "Scope1" -StartRange 192.168.1.100 -EndRange 192.168.1.200 -SubnetMask 255.255.255.0

- Use PowerShell to create a new superscope named Superscope1 by taking the existing scopes Scope1 (192.168.1.100 to 192.168.1.200) and Scope2 (192.168.2.100 to 192.168.2.200):

  Add-DhcpServerv4Superscope -SuperscopeName "Superscope1" -ScopeId 192.168.1.0, 192.168.2.0

- Use the Netsh command to add a new scope named Scope1 for the 192.168.1.0/24 network and then add an IP address range from 192.168.1.100 to 192.168.1.200 to the scope (note that it is two commands):
  1. **netsh dhcp server add scope 192.168.1.0 255.255.255.0 “Scope1”**
  2. **netsh dhcp server scope 192.168.1.0 add iprange 192.168.1.100 192.168.1.200**

Configure a DHCP reservation

For the exam, know how to add a reservation by using the DHCP management console, PowerShell, and Netsh. Below are some real world examples.

- Use PowerShell to add a reservation for Project #14 by reserving 192.168.1.75:

Add-DhcpServerv4Reservation -ScopeId 192.168.1.0 -IPAddress 192.168.1.75 -ClientId F1-5D-4B-00-3C-82 -Description "Projector #14"

- Use Netsh to add a reservation to reserve 192.168.1.75:

  netsh dhcp server scope 192.168.1.0 add reservedip 192.168.1.75 F15D4B003C82

To add a reservation using any of the available methods, you need a MAC address and the IP address that you want to reserve.

Note: you can have a reservation for an IP address that is part of an excluded range.

Configure DHCP options

DHCP options consist of two areas: server options and scope options. Server options are server wide and are part of every scope. Scope options are scope wide and only impact the scope.

Below is a screen capture of the server and scope options window showing some of the available options. Server and scope options are configured in the same way, albeit each has its own configuration.

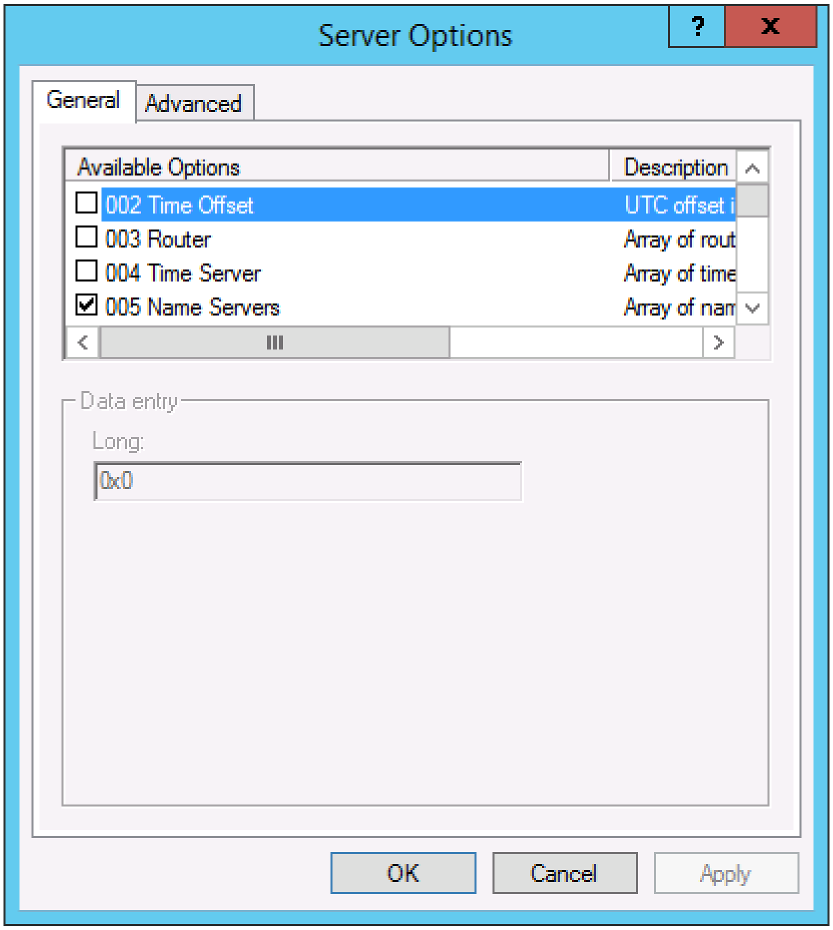
Server Options

Available Options	Description
□ 002 Time Offset	UTC offset i
□ 003 Router	Array of rout
□ 004 Time Server	Array of time
☑ 005 Name Servers	Array of nam

Data entry

Long:

0x0

Below is a screen capture of some defined scope options.

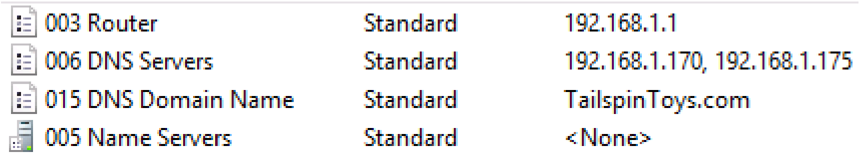

003 Router	Standard	192.168.1.1
006 DNS Servers	Standard	192.168.1.170, 192.168.1.175
015 DNS Domain Name	Standard	TailspinToys.com
005 Name Servers	Standard	<None>

For the exam, know the following information:

- **Scope options take precedence over server options**. If you configure a setting in the server options and configure it different in the scope options, clients will get the settings based on the scope options.
- **There are class options**. Class options apply to clients that specify a specific DHCP class ID. Class options take precedence over scope and server options. You can modify the class options by using the Advanced tab in the scope or server options. Note that, by default, the options are quite limited

and mostly for legacy systems such as Windows 2000. However, third-party vendors can extend class options.

- **There are reserved client options**. These are used for clients with reservations. These take precedence over all other settings. To use these, you first must create a DHCP reservation. Thereafter, you can configure options directly on the reservation.

Configure client and server for PXE boot

Configuring PXE boot is a fairly straight forward process – set up the PXE environment, such as Windows Deployment Services (WDS), ensure that DHCP is functional, and boot the client to PXE. However, there are some important details that you should be aware of for the exam, including the following:

- **Use DHCP forwarders when clients are on a different subnet**. If your clients are on one subnet and your server infrastructure (DHCP and PXE) are on another subnet, you need to have a solution to get the DHCP broadcast packets from the client subnet to the server subnet. DHCP forwarders, sometimes called IP helpers or DHCP relay agents, can be used to forward DHCP requests from one subnet to the DHCP server. Watch for exam scenarios that call out multiple subnets with PXE.
- **Configure DHCP options if you can't use forwarders**. If your DHCP server and WDS server are different servers and on different subnets, you can configure DHCP options to enable clients to boot to find the WDS server and boot to PXE. You need to configure Option 60 with a value of PXEClient. You need to configure Option 66 with the hostname or IP address of the WDS server. You need to configure Option 67 with the path to the boot file name (boot\x86\wdsnbp.com).
- **DHCP forwarders / IP helpers are the preferred method**. Microsoft recommends to use DHCP forwarders / IP helpers when possible because they provide a more robust solution for clients booting to PXE. Some of the drawbacks of using DHCP options are that the WDS server is hardcoded (so if you have multiple servers, you can't take advantage of load balancing), WDS offers multiple boot files and hard coding the boot file may reduce performance for some clients, such as 64-bit clients.

Configure DHCP relay agent

You can use any DHCP forwarder in your environment. However, for the exam, you should be familiar with adding and configuring a DHCP forwarder on Windows Server 2012 R2. Here are the key pieces of information to know:

- **To add a DHCP forwarder, you first have to add the Remote Access role and the DirectAccess and VPN (RAS) role service**. Back in Windows Server 2003, a DHCP forwarder was added by adding a new routing protocol in the Routing and Remote Access management console. This leads to a potentially confusing situation for the exam. There is a role service named Routing under the Remote Access role. It may look compelling on the exam. But it isn't needed for adding a DHCP forwarder.
- **There are two things that have to be set for DHCP forwarding to work**. They are:

- **An interface (NIC) has to be enabled for relay mode**. And, if the server is multi-homed, you need to choose the interface that is on the same subnet as the DHCP clients. You can enable relay mode on multiple interfaces but this would be a fairly uncommon configuration. Watch for exam situations where a DHCP relay agent is already enabled for relay mode but DHCP relay isn't working. In such a scenario, you should add the IP address of the DHCP server to the configuration or ensure that the correct interface is used for DHCP forwarding.
- **The IP address of the destination DHCP server must be added to the properties of the DHCP Relay Agent node in the Routing and Remote Access management console**. Note that you can add more than one DHCP server. In such a case, the first DHCP server that responds to the client will become the DHCP server for the client.

The screen captures below show the Routing and Remote Access management console and the location of the DHCP forwarding configuration.

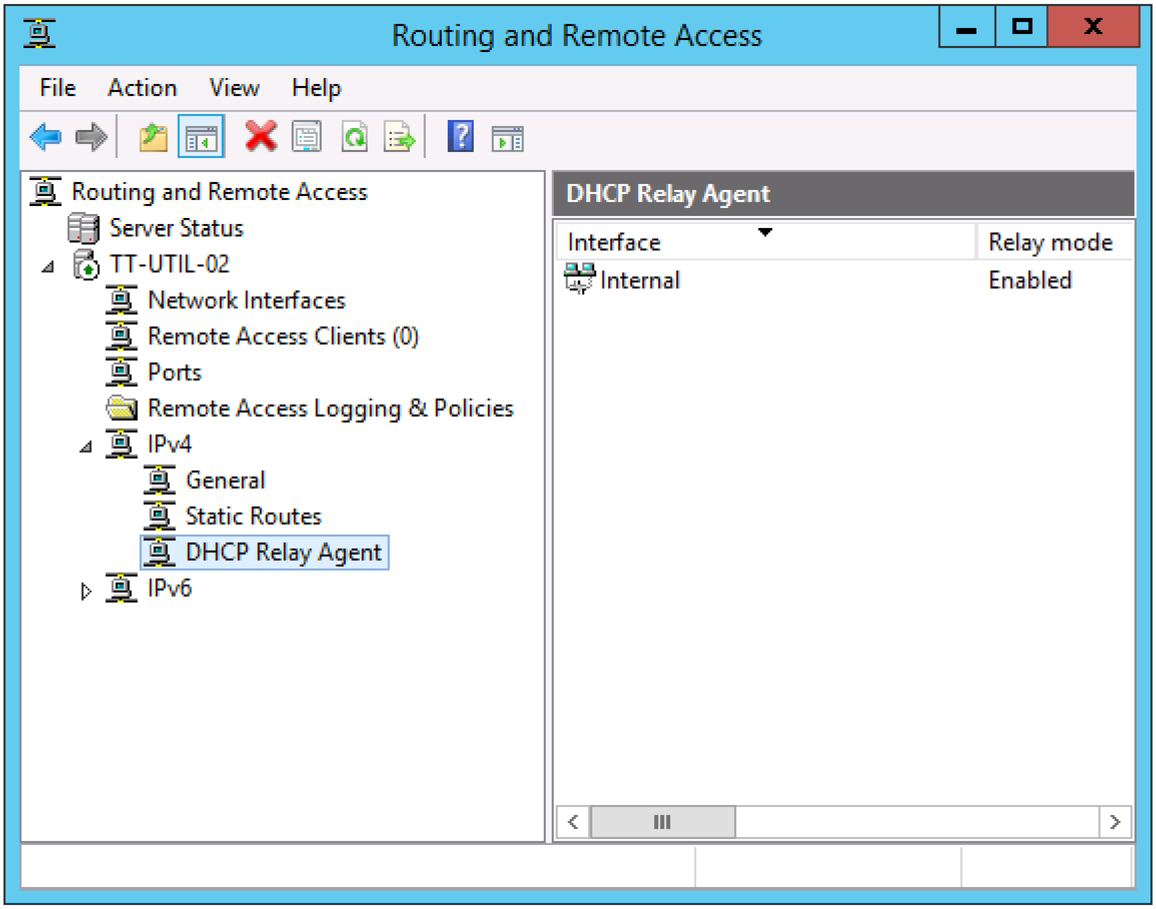

Interface	Relay mode
Internal	Enabled

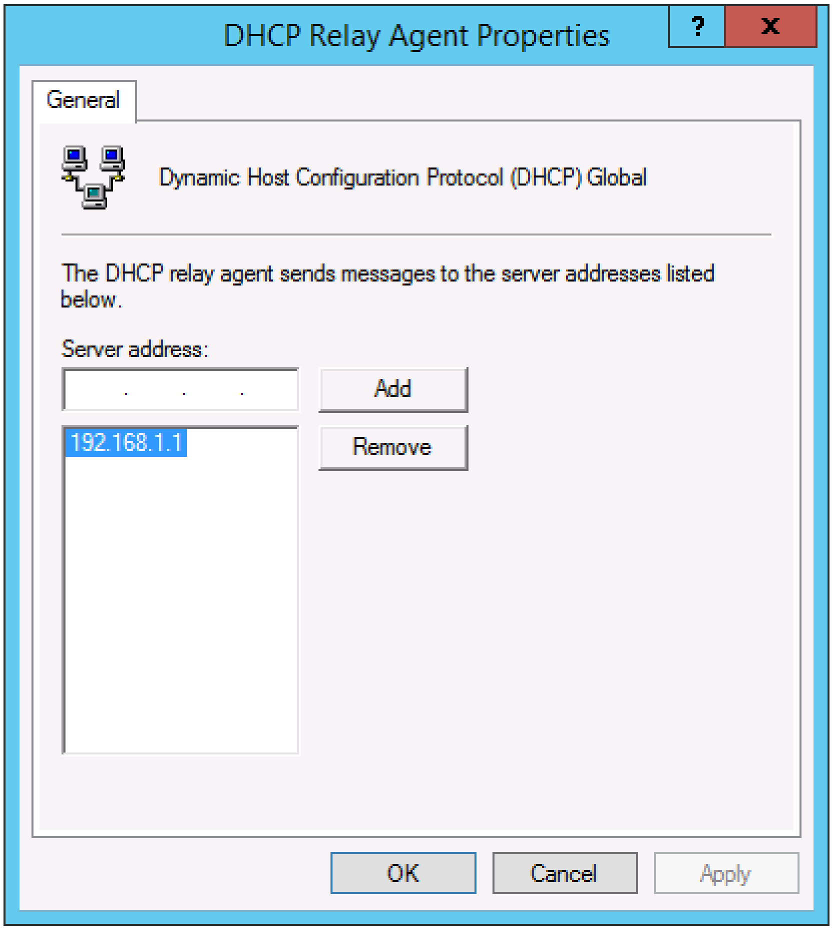

DHCP Relay Agent Properties

Dynamic Host Configuration Protocol (DHCP) Global

The DHCP relay agent sends messages to the server addresses listed below.

Server address: 192.168.1.1

Authorize DHCP server

In order for a domain joined DHCP server to issue IP addresses to clients, the DHCP server must be authorized in Active Directory. Standalone servers can issue IP addresses to clients that are joined to a domain and do not have to be authorized. Be sure to carefully read exam questions so you do not miss details that are key to answering the questions. For the exam, know the following information:

- **If you add DHCP to a domain controller, it is automatically authorized**. However, a good practice is to run DHCP on a member server, not a domain controller.
- **To authorize a DHCP server on a member server, you need to be a member of the Enterprise Admins group**. Otherwise, you have to be delegated rights specifically to authorize a DHCP server. The exam likely will not get into the delegated rights. You can delegate rights by using Active Directory Sites and Services and delegating at the Services/NetServices node.

Deploy and configure DNS service

DNS is one of the foundational technologies of a network. Most administrators are familiar with DNS at a high level. For the exam, you also need to understand the integration with Active Directory and how to configure DNS by using PowerShell and the command line.

Configure Active Directory integration of primary zones

DNS zones that are not Active Directory integrated are stored as text files on a DNS server. DNS servers synchronize DNS changes by using zone replication. But for Active Directory integrated zones, DNS information is stored in the Active Directory database (NTDS.DIT) and DNS synchronization occurs by way of Active Directory replication. By default, new zones created in the DNS management console are Active Directory integrated. The screen capture below shows the default setting which dictates that a zone will be Active Directory integrated.

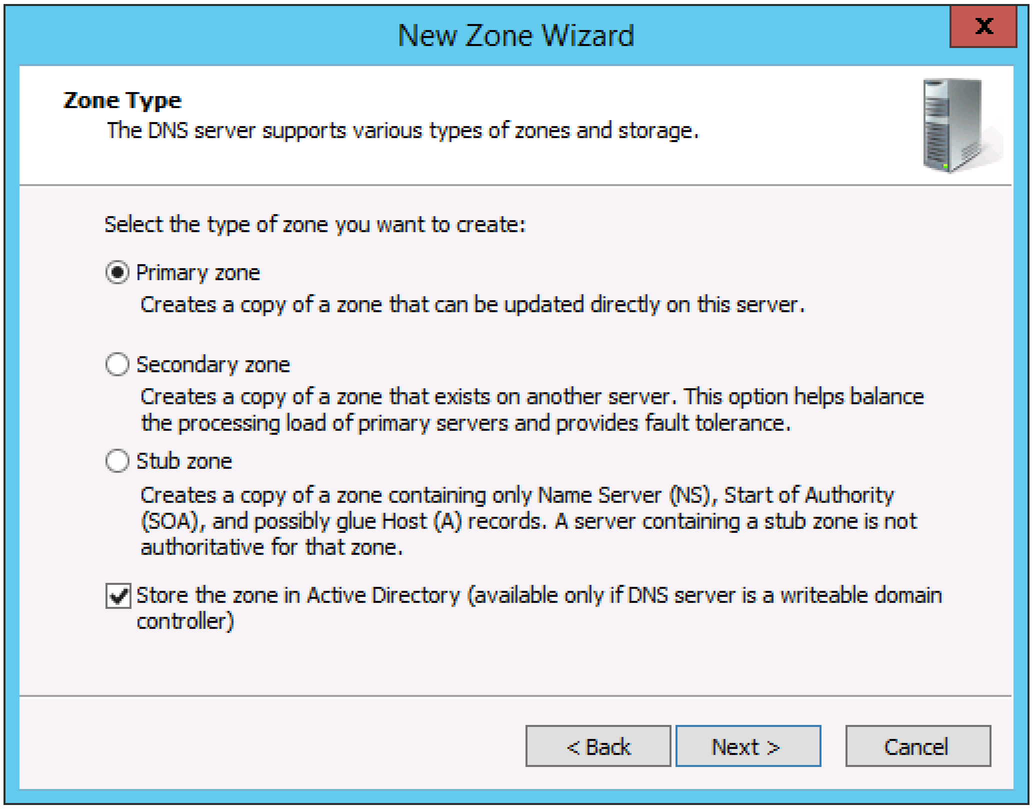

New Zone Wizard

Zone Type
The DNS server supports various types of zones and storage.

Select the type of zone you want to create:

● Primary zone
Creates a copy of a zone that can be updated directly on this server.

○ Secondary zone
Creates a copy of a zone that exists on another server. This option helps balance the processing load of primary servers and provides fault tolerance.

○ Stub zone
Creates a copy of a zone containing only Name Server (NS), Start of Authority (SOA), and possibly glue Host (A) records. A server containing a stub zone is not authoritative for that zone.

☑ Store the zone in Active Directory (available only if DNS server is a writeable domain controller)

When you create a new DNS zone with PowerShell, you specify whether a zone is Active Directory integrated by specifying the replication scope, as shown in the examples below.

- Create new AD integrated zone named tailspintoys.com with replication to all domain controllers in the domain that run DNS:

  Add-DnsServerPrimaryZone -Name “tailspintoys.com” -ReplicationScope “Domain”

- Create a new AD integrated zone named tailspintoys.com with replication to all domain controllers in the forest that run DNS:

  Add-DnsServerPrimaryZone -Name “tailspintoys.com” -ReplicationScope “Forest”

- Create a new zone named tailspintoys.com that is not Active Directory integrated:

  Add-DnsServerPrimaryZone -Name “tailspintoys.com” -ZoneFile “tailspintoys.com.dns”

You can also create new zones by using the dnscmd.exe command line tool. To create an Active Directory integrated zone for a domain named tailspintoys.com, you can run the **dnscmd /zoneadd tailspintoys.com /dsprimary** command. The key to the command is the /dsprimary switch. For a non-integrated zone, you use the /primary switch instead.

Configure forwarders

For this part of the exam, you should understand server forwarders. You may find some content related to conditional forwarders too. Types of forwarders:

- **Server forwarders**. A server forwarder is configured at the server level. A server forwarder is often used to resolve names that the server doesn’t know about. In some companies, administrators configure their DNS servers to forward to their ISP’s DNS servers. A server forwarder is used for all DNS queries that cannot be resolved by the server.
- **Conditional forwarders**. A conditional forwarder is configured under the Conditional Forwarders node in the DNS management console. Conditional forwarders can be stored in Active Directory and replicate by using Active Directory replication. Conditional forwarders are used to forward DNS queries for a single specified domain to a specified DNS server or set of DNS servers. The primary difference between a server forwarder and a conditional forwarder is that a server forwarder is not domain specific while a conditional forwarder is domain-specific.

For the exam, know the following information:

- **Conditional forwarders take precedence**. Imagine that you are part of the contoso.com domain and you set up a conditional forwarder for tailspintoys.com. Then, you add a server forwarder pointing to your ISP’s DNS server. When you try to resolve a host in the tailspintoys.com domain, the conditional forwarder will be used.

- **Root hints can take over if forwarders are not functional**. There is an option to enable the use of root hints if configured forwarders are not functional.
- **Server forwarders are server specific**. You must individually configure servers for forwarders. Thus, you could have one server forwarding and another not forwarding.
- **Using PowerShell to create forwarders**. You can use the Add-DnsServerConditionalForwarderZone PowerShell command to add a conditional forwarder. You can use the Add-DnsServerForwarder command to add a server forwarder.

Configure Root Hints

Root hints are configured on a per server basis. By default, Windows includes the root hints and they are available for use after installing the DNS Server role. For the exam, know the following information:

- **Root hints are not always used**. Root hints are used only if server forwarders are not being used or if server forwarders are not responding and your server is configured to use root hints when server forwarders do not respond.
- **Root hint servers occasionally need to be updated because a server is changed or an IP address changes**. You can download the latest root hints files at https://www.iana.org/domains/root/files.
- **There are five PowerShell commands that you can use to manage root hints**. The commands are:
  - Add-DnsServerRootHint - adds a root hint to the existing root hints list
  - Get-DnsServerRootHint - displays the existing root hints
  - Import-DnsServerRootHint - copies root hints from one server to another but does not overwrite existing root hints
  - Remove-DnsServerRootHint - removes one or more root hints from the server
  - Set-DnsServerRootHint - overwrites existing root hints with the root hints you specify with this command

Manage DNS cache

From a server perspective, there aren't a lot of DNS cache management tasks. You should understand the following information for the exam:

- **You need to use the advanced view to see the cache**. In the DNS management console, you can use the advanced view to display the Cached Lookups node.
- **You can individually delete cached items from the cache**. This is helpful when you have a bad DNS record but don't want to clear the entire cache.
- **You can delete the entire DNS cache on a server**. To do so, right-click the DNS server and then click **Clear Cache**. This is helpful in a scenario where a large amount of the cached entries are invalid.

- **You can use PowerShell to perform cache-related management tasks**. Know the following commands:
    - Show-DnsServerCache displays all of the cached DNS entries on a server.
    - Clear-DnsServerCache -Force deletes all of the cached DNS entries on a server.
    - Get-DnsServerCache displays the current cache configuration settings such as the maximum TTL (dictates the maximum time an entry will be cached) and the locking percent. The locking percent dictates the point at which a cached entry can be overwritten. Other parameters that you can view (and set with another command) are the MaxKBSize (the maximum amount of storage to use for caching), the MaxNegativeTtl (the amount of time to cache a negative answer for a DNS query), PollutionProtection (dictates whether filtering is used for name service DNS records), and StoreEmptyAuthenticationResponse (dictates whether empty responses from authoritative servers are stored in the cache).

Create A and PTR resource records

For the exam, you should know how to create A records and PTR records using the DNS management console. Because those are routine administrative tasks and most administrators are familiar with them, I will only cover the PowerShell and command line methods that you should know:

- Use PowerShell to create a new A record for a server named Server 19 with an IP address of 192.168.254.200 in the tailspintoys.com domain:

  Add-DnsServerResourceRecord -ZoneName tailspintoys.com -A server19 -IPv4Address 192.168.254.200

- Use PowerShell to create a new PTR record for 192.168.254.200 pointing to server19.tailspintoys.com:

  Add-DnsServerResourceRecord -ZoneName “254.168.192.in-addr.arpa” -PtrDomainName “server19.tailspintoys.com”

- Use dnscmd.exe to create a new A record for a server named Server 19 with an IP address of 192.168.254.200 in the tailspintoys.com domain:

  dnscmd /recordadd server19 A 192.168.254.200

- Use dnscmd.exe to create a new PTR record for 192.168.254.200 pointing to server19.tailspintoys.com:

  dnscmd /recordadd 254.168.192.in-addr.arpa 200 PTR server19.tailspintoys.com

Be aware that dnscmd.exe is deprecated. Thus, it may not be available in future versions of Windows. Generally, exam item writers avoid deprecated technologies unless they are called out specifically in the exam OD. But it is still helpful to know the dnscmd.exe syntax just in case.

5 - Install and administer Active Directory (14%)

At first glance, it appears that Active Directory only represents a small portion of the 70-410 exam. However, this published percentage is a bit misleading. Only the topics listed below make up the 14%. When somebody describes Active Directory, they virtually always include Group Policy and DNS. Group Policy has its own functional group on the exam and DNS has a dedicated section in the previous functional group. So when you look at the big picture, Active Directory and the technologies that make it work or enhance it, make up a big chunk of the exam.

Install domain controllers

This section on installing domain controllers focuses on installation types that many administrators aren't working with regularly. I recommend that you walk through the installation of a domain controller running on a Server Core installation in a lab and also try installing a domain controller from media. These two walk-throughs in a lab would be a big help in remember the information for the exam.

Add or remove a domain controller from a domain

The DC promotion process has changed since Windows Server 2012 was released. If you haven't promoted a DC on Windows Server 2012 or 2012 R2, you should do so prior to taking the exam. You should know the following information about adding domain controllers to the domain:

- **You cannot promote a domain controller until the Active Directory Domain Services role is installed**. Thus, watch for scenarios where you are trying to promote a server to a domain controller but the role hasn't been added yet.
- **In Server Manager, after the Active Directory Domain Services role has been added, the Notifications section will display a warning icon indicating that there are post-deployment tasks that are required**. You can click to view the notification and you'll see the following notification:

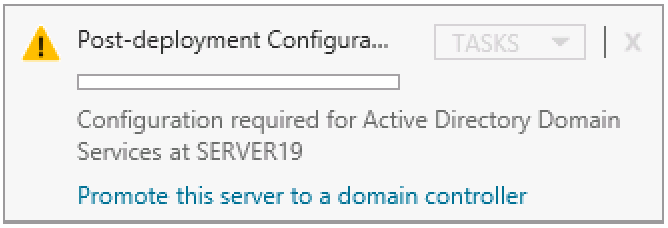

Post-deployment Configura...

Configuration required for Active Directory Domain Services at SERVER19

Promote this server to a domain controller

You can click the **Promote this server to a domain controller** link to start the promotion wizard.

- **You can use PowerShell to add roles**. To add the Active Directory Domain Services role by using PowerShell, you run the **Add-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools** command.
- **You can use PowerShell to promote a server**. To promote a server to a domain controller by using PowerShell while installing the DNS role in an existing domain named tailspintoys.com, you run the **Install-ADDSDomainController -InstallDns -Credential (Get-Credential tailspintoys\administrator) -DomainName tailspintoys.com** command. You will be prompted for the Directory Services Restore Mode (DSRM) password.
- **You can demote a domain controller by using PowerShell's Uninstall-ADDSDomainController cmdlet**. After the demotion, you should remove the Active Directory Domain Services role.
- **To demote a domain controller by using the GUI, you can use Server Manager to remove the Active Directory Domain Services role**. During the role removal, you will be prompted to demote the domain controller.
- **The Dcpromo tool does not work for regular promotions and demotions since Windows Server 2012**. However, you can still use it for unattended installations.

Install Active Directory Domain Services (AD DS) on a Server Core installation

The process of adding a domain controller on a Server Core installation is the same as a standard installation if you use PowerShell. First, you add the Active Directory Domain Services role. Then you use the Install-ADDSDomainController cmdlet to promote the server. You can also use Dcpromo with an answer file. You should be familiar with an unattended installation.

The command to use for an unattended installation is **dcpromo /unattend:C:\temp\unattend.txt**.

Below is a sample unattend.txt file for a domain named alpineskihouse.com:

```
UserName=Administrator

UserDomain=alpineskihouse.com

Password=
ReplicaDomainDNSName=alpineskihouse.com

ReplicaOrNewDomain=Replica

DatabasePath=D:\DB\NTDS

LogPath=E:\Logs\NTDS

SYSVOLPath="%systemroot%\SYSVOL"

InstallDNS=Yes

ConfirmGC=Yes

SafeModeAdminPassword=

RebootOnCompletion=Yes
```

Install a domain controller from Install from Media (IFM)

When you have a backup of a domain controller, you can use it to promote a new domain controller. In such a scenario, the new DC only has to replicate the changes since the backup, not the entire Active Directory database.

Perform the following steps:

1. From a DC, run the ntdsutil program.
2. From ntdsutil, run the following commands:
   - **activate instance ntds**
   - **ifm**
   - **create sysvol full D:\backups**
3. Copy the data from the backup to the new server that will be promoted.
4. Add the Active Directory Domain Services role to the server that will be promoted.
5. Start the domain controller promotion process.
6. Select the **Add a domain controller to an existing domain** option. Specify the domain name and credentials, if necessary and then click **Next**.
7. Select the desired domain controller options, type a directory services restore mode password, and then click **Next**.
8. Select the desired options and the **Install from Media** option. Type the path to the IFM files and then specify the replication from a specific domain controller if desired. Then click **Next**.
9. Specify the database, log, and SYSVOL paths and then click **Next**.
10. Review the summary of options and then click **Next**.

11. Click **Install** to complete the promotion.

Resolve DNS SRV record registration issues

SRV records are crucial to the functionality of domain controllers. You need to be familiar with the following information for the exam:

- **Where to find the SRV records that get registered**. When you promote a server to a domain controller, all of the SRV records that are registered are stored in %systemroot%\system32\config\netlogon.dns.
- **How to force SRV record registration**. To force a domain controller to reregister SRV records, you can perform the following tasks:
  - Run the **net stop netlogon** command and then the **net start logon** command.
  - Run the **Restart-Service netlogon** PowerShell command.
  - Manually create the SRV records found in the netlogon.dns file.
  - Run the **dcdiag /fix** or **netdiag /fix** commands. Note that the netdiag tool is not available as part of a default installation of Windows Server.
- **How to troubleshoot SRV registration issues**. If a domain controller is not registering its SRV records, check the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** registry location for an entry named **DnsAvoidRegisterRecords**. Such an entry may be configured to prevent a domain controller from registering some DNS records. This configuration is sometimes used to prevent authentication or use of other services for specific domain controllers due to firewall or other restrictions.
- **Optional settings that control how SRV records are used**. There are two other important optional registry values for SRV records. One is **LdapSrvPriority** which controls the order that domain controllers are contacted in (a value of 0x0 is the highest priority and is the default value while a value of 0xFFFF is the lowest priority). The other is **LdapSrvWeight** which determines the order that domain controllers are contacted in if their priorities are the same. The default value is 0x64 (which represents 100). By increasing that number, a domain controller can be configured for more weight (and thus a higher chance of being the chosen domain controller during communication).

Configure a global catalog server

Configuring a domain controller to be a global catalog server is very straight forward in the GUI. Open Active Directory Sites and Services, navigate to the NTDS Settings object of a domain controller, right-click **NTDS Settings** and bring up the **Properties** page. From there, you can click the Global Catalog option. However, for the exam, you should know how to do this outside of the GUI too:

- **Use PowerShell to configure a domain controller as a global catalog server**. This isn't as straight forward as many of the administrative tasks such as creating users or groups. To configure DC1 in the alpineskihouse.com domain as a global catalog server, you run the following command:

  ```
  Set-ADObject "cn=NTDS
  Settings,cn=DC1,cn=Servers,cn=Site1,cn=Sites,cn=Configuration,dc=alpineskihouse,dc=com" -
  Replace @{Options='1'}
  ```
- **Use the dsmod command-line utility to configure a domain controller as a global catalog server**. Run the following command:

  ```
  Dsmod server "cn=NTDS
  Settings,cn=DC1,cn=Servers,cn=Site1,cn=Sites,cn=Configuration,dc=alpineskihouse,dc=com" -
  isgc Yes
  ```

You should also understand the following information about global catalog servers:

- **In a single domain forest, the global catalog and a standard domain controller perform the same tasks**. Thus, it is a good practice to configure all domain controllers as global catalog servers.
- **In a multi-domain environment, global catalog servers require a little extra disk space to store a partial replica of other domains in the forest**. It also uses additional bandwidth to replicate that data.
- **After you configure a global catalog server in a multi-domain environment, replication is required to get the global catalog ready to service global catalog requests**. Once it is ready, Event ID 1119 will be logged in the Directory Services event log which indicates that the server is servicing global catalog requests.
- **The global catalog listens on TCP port 3268**. The port must be opened if communication goes through a firewall.

Deploy Active Directory infrastructure as a service (IaaS) in Microsoft Azure

You can deploy AD DS in Azure in much the same way as you do on-premises. The process of promoting a domain controller is the same. The process of adding sites and subnets is the same. The management tools and processes are the same. But there are some key differences and you should be familiar with them:

- **You need to establish connectivity between your environment and Azure**. Azure provides connectivity by way of their Virtual Network solution.
- **You should assign a static IP to your Azure-based DCs**. By default, VMs in Azure have a dynamic IP address. Instead, you should set a static IP address to maximize the uptime and functionality of your AD DS environment. Note that you have to set a static IP by using PowerShell's Set-AzureStaticVNetIP cmdlet.

You should be familiar with Azure's Virtual Network solution which offers 3 ways to connect your environment to Azure:

- **Point-to-site**. A point-to-site VPN connects a specified computer to Azure over a VPN. You can configure multiple computers to use a point-to-site VPN simultaneously. This type of solution is best for early testing and one-off nonproduction uses. Routing is static.
- **Site-to-site**. A site-to-site VPN connects your network to an Azure VPN Gateway. This solution is best for nonproduction workloads and small production workloads. Routing is static.
- **ExpressRoute**. ExpressRoute is Azure's premier connectivity solution. When you connect your networks to Azure by using ExpressRoute, traffic does not go over the public Internet. It offers the highest performance and security of any of the connectivity options. To use ExpressRoute, you choose a provider that offers ExpressRoute circuits and you can connect the circuit to an existing MPLS network. You can also connect ExpressRoute to a third-party co-location facility.

Create and manage Active Directory users and computers

Be prepared to know the PowerShell method for all of the tasks in this section of the exam. Outside of PowerShell, the exam item writers will likely target some of the newer GUI tools such as the Active Directory Administrative Center over established tools such as Active Directory Users and Computers.

Automate the creation of Active Directory accounts

When you think about automating the creation of user accounts, you might be thinking about templates. But templates are covered later in this section. For automating the creation, be thinking about PowerShell and other methods outside of the GUI.

You should know how to use the ldifde utility to import users. First, you need an import file (an .ldf file) with the following syntax for each user:

```
dn: cn=Brian Svidergol,ou=Engineering,dc=adatum,dc=com

changetype: add

cn: Brian Svidergol

objectClass: user

samAccountName: bsvidergol

mail: brian@svidergol.com
```

You can create import files by exporting users using ldifde. Or, you can create the import file other ways. Once the file is ready, you can run the following command to create users from the import file:

```
Ldifde -i -f users.ldf
```

If you need to create a large amount of test users, you can automate the creation of them. For example, to create 1,000 users in the Test OU in the adatum.com domain with a naming convention of User1, User2, User3 (and so on), run the following command:

```
for /L %i in (1,1,1000) do dsadd user "cn=User%i,ou=Test,dc=adatum,dc=com" -samid User%i -pwd N7dd__WECg8ie.%i
```

You can use PowerShell to create 1,000 test users too, as follows:

```
$u=1

Do {New-ADUser -Name User$u

$u++

} while ($u -le 1000)
```

You can also use PowerShell to import users from a .CSV file by using the Import-Csv cmdlet and piping the imported .csv file to a ForEach.

Create, copy, configure, and delete users and computers

This section is about object management. The tasks are common and often performed multiple times per day by administrators. For the exam, study the PowerShell methods and command-line methods. If you don't work with Active Directory on a regular basis, you should take a quick look at the GUI to ensure that you are familiar with the GUI method too.

- **Create a new user with PowerShell**. Run the following command:

  ```
  New-ADUser -Name “User1” -Path “ou=Employees,ou=HQ,dc=adatum,dc=com” -GivenName “User1” -sAMAccountName “user1”
  ```

- **Copy User1 to User2**. In this case, use the GUI. Browse to a user in Active Directory Users and Computers, right-click the user object, and then click Copy to begin the process.

- **Configure an attribute for a user**. Use PowerShell's Set-ADUser cmdlet, as follows:

  ```
  Set-ADUser -Identity "cn=User1,ou=Employees,ou=HQ,dc=adatum,dc=com" -Description
  "Added to Group1 on 8/7/15"
  ```

- **Configure the home directory path for all users in the Employees OU**. Run the following PowerShell command:

  ```
  Get-ADUser -SearchBase "ou=Employees,ou=HQ,dc=adatum,dc=com" -Filter * | Set-ADUser -
  HomeDirectory "\\Server2\%username%"
  ```

- **Delete a user**. Run the following PowerShell command. Notice that the confirmation is false which means that you will not get prompted to confirm the deletion.

  ```
  Remove-ADUser -Identity "cn=User1,ou=Employees,ou=HQ,dc=adatum,dc=com" -
  Confirm:$False
  ```

- **Delete an OU that contains multiple user objects**. Run the following PowerShell command to delete all users in the Employees OU without confirmation:

  ```
  Remove-ADObject -Identity "ou=Employees,ou=HQ,dc=adatum,dc=com" -Recursive -
  Confirm:$False
  ```

To prepare for the exam, you should run through some of these commands in a lab environment. By reading them and then running them, your retention for the exam will be higher.

Configure templates

You can create a user object in Active Directory and use it as a template. Once created, you can use the GUI tools to copy it and create new users. Copying it retains some of the attributes for the newly created user object. The attributes that are copied are marked in the schema and you can configure your environment to copy additional attributes, if desired. You won't need this information for the exam, but if you are interested to learn more, have a look at Recipe 10.11 in the Active Directory Cookbook, Fourth Edition.

You can use PowerShell to create a new user from a template by using a Get-ADUser command and a New-ADUser command. The following example is two lines:

```
$user = Get-ADUser -Identity "cn=Template1,ou=Misc,ou=HQ,dc=adatum,dc=com" -Properties
department, co, title, l, c, st, countrycode
```

New-ADUser -Instance $user -Name "User1" -DisplayName "User1" -GivenName "User1" -UserPrincipalName "User1@adatum.com" -SamAccountName "User1" -PasswordNotRequired $true -Enable $true

Perform bulk Active Directory operations

You can perform some bulk operations in the GUI. In Active Directory Users and Computers, multi-select several users, right-click a highlighted user, and then click Properties. The Properties window that is displayed is the multi-user window which enables you to modify attributes for all of the selected users. You can change the following fields: Description, Office, Telephone number, Fax, Web page, E-mail, UPN suffix, Logon hours, Computer restrictions, any of the account options such as requiring a password change at next logon, account expiration, all of the address fields, all of the profile fields, and all of the organization fields. Attributes that must have a unique value, such as the logon name, cannot be changed during a multi-select. You can also disable and enable user accounts with multi-select. One lesser known fact is that multi-select is not limited to a single container in Active Directory Users and Computers. However, if you try to find users and then multi-select to update fields, it does not work. Instead, you have to create a saved query and then you can multi-select objects from the query without regard for their location.

With PowerShell, you can run a Get-ADUser command to target users by location or filter and then pipe the output to a Set-ADUser command, as shown previously. You should be familiar with the attributes that you can modify with Set-ADUser. See the Set-ADUser command reference for a complete list at https://technet.microsoft.com/en-us/library/ee617215.aspx.

Configure user rights

For this section of the exam, user rights is not the user rights assignments that you control by using Group Policy. Those are covered in the functional group titled “Create and manage Group Policy”. For this section, user rights will mean group membership, delegation, administrative access, and tasks related to giving users rights to perform specific actions. You should be familiar with the following tasks:

- **Add User1 to Group1 using PowerShell**. Run the following command:

  Add-ADGroupMember -Identity “cn=Group1,ou=Groups,dc=adatum,dc=com” -Members “cn=User1,ou=Employees,dc=adatum,dc=com”

- **Remove User1 from Group1 using PowerShell**. Run the following command:

  ```
  Remove-ADGroupMember -Identity "cn=Group1,ou=Groups,dc=adatum,dc=com" -Members "cn=User1,ou=Employees,dc=adatum,dc=com" -Confirm:$False
  ```

Offline domain join

Offline domain join is a technology that enables computers to join a domain without contacting a domain controller during the joining process. For the exam, you should be familiar with the overall process from start to finish. Know the following information:

- **Client computers must run Windows 7 or later**.
- **Server computers must run Windows Server 2008 R2 or later**.
- **To perform an offline domain join, you must have rights to join computers to the domain**.
- **Using offline domain join is a two-step process**. Perform the following steps:
  - **Provision the computer**. Run the **djoin /provision /domain adatum.com /machine Computer1 /savefile djoin.txt** command. You can run this command from a domain controller or domain member computer.
  - **Put the computer account metadata on the computer that you are going to join offline**. Run the **djoin /requestodj /loadfile djoin.txt /windowspath %windir% /localos** command. You need to run this command on the computer that you are joining to the domain.
- **You can use offline domain join for computers that are online and running Windows**. Or, you can use the offline domain join process during the unattended installation of Windows. For unattended installations, you need to modify the Unattend.xml file with a section for the offline domain join.
- **For the exam, be aware of when offline domain join might be suitable**. The following are two examples:
  - **Reduce the time required to deploy computers (or VMs)**. Being able to deploy new VMs that are already joined to the domain reduces the total time required to deploy new computers.
  - **Provide domain joining capabilities to computers that do not have connectivity to a domain controller**. For example, if you were preparing a new branch office for opening, you may not have network connectivity and firewalls established to join computers to the domain. But with offline domain join, you could join all of the computers to the domain prior to the network connectivity and firewall configuration.

Manage inactive and disabled accounts

You need to know how to find user an computer objects that are inactive (often referred to as “stale”). The following PowerShell code can be run from the PowerShell prompt or saved as a script. It finds users that have not logged on in the last 90 days and then outputs the account name and the last date of logon:

```
$DaysSince = (Get-Date).AddDays(-90)
```

```
Get-ADUser -Filter * -Properties LastLogonDate | Where-Object {($_.LastLogonDate -le $DaysSince)  -and ($_.Enabled -eq $True) -and ($_.LastLogonDate -ne $NULL)} | Select Name,LastLogonDate
```

You can use the following command to find computers that have not logged on in the last 90 days:

```
$DaysSince = (Get-Date).AddDays(-90)
```

```
Get-ADComputer -Filter * -Properties LastLogonDate | Where-Object {($_.LastLogonDate -le $DaysSince) -and ($_.Enabled -eq $True) -and ($_.LastLogonDate -ne $NULL)} | Select Name,LastLogonDate
```

To find all disabled users in the domain, run the following PowerShell command:

```
Get-ADUser -Filter {Enabled -eq “False”} | FL Name
```

To find all disabled computers in the domain, run the following PowerShell command:

```
Get-ADComputer -Filter {Enabled -eq “False”} | FL Name
```

You can also find disabled users or users that haven’t logged on recently by using Active Directory Users and Computers. In the screen capture below, the query will return disabled users:

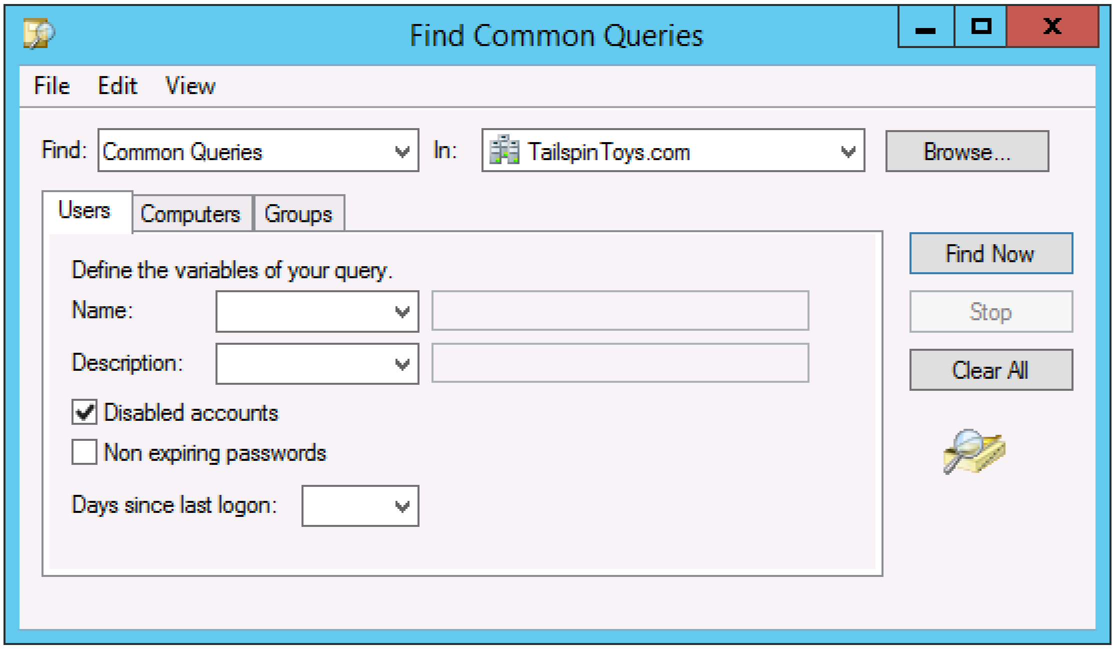

Find: Common Queries In: TailspinToys.com

Define the variables of your query.

Name:

Description:

☑ Disabled accounts

□ Non expiring passwords

Days since last logon:

You can also use Active Directory Users and Computers to find disabled computer objects, as shown in the screen capture below:

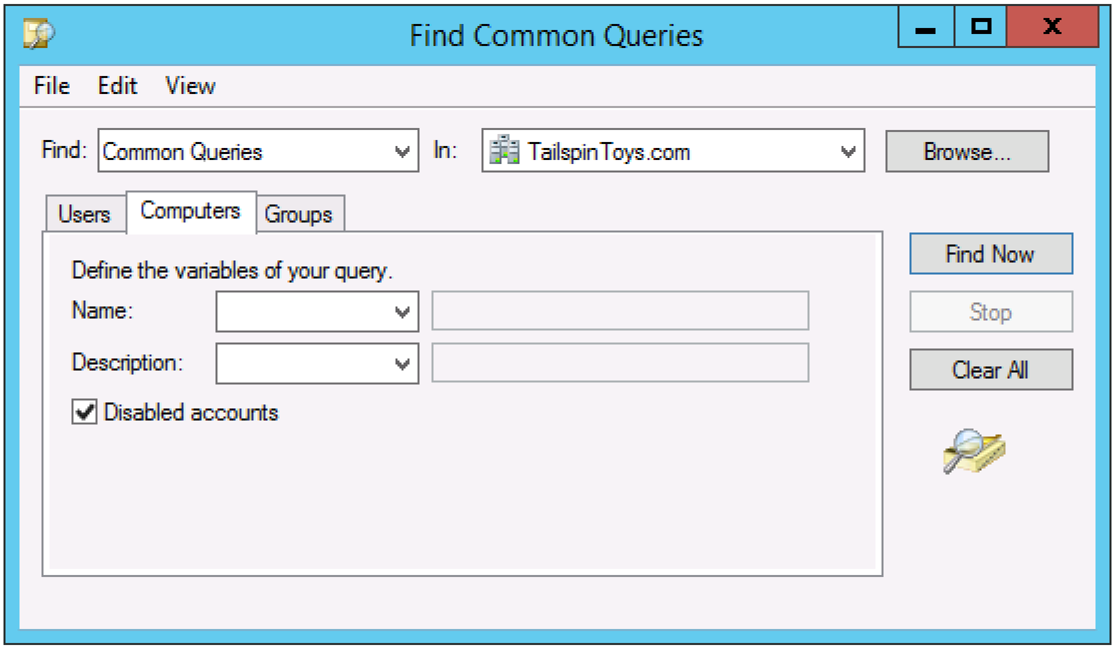

Find: Common Queries In: TailspinToys.com

Define the variables of your query.

Name:

Description:

☑ Disabled accounts

Create and manage Active Directory groups and organizational units (OUs)

This section is dedicated to groups and OUs. Creating and managing groups and OUs using the GUI tools is very straight forward. In fact, many of the tasks have been the same in the GUI tools for the last 15 years. Thus, expect that the exam item writers have looked for other ways to test your knowledge and skills in these areas. I suspect that PowerShell and other command-line methods will be prevalent for this section.

Configure group nesting

Group nesting occurs when one group is a member of another group. Nesting was originally considered one way to reduce Active Directory replication traffic. The downside of nesting, especially with several levels of nesting, is that it makes troubleshooting more difficult. For the exam, you should be aware of the limitations of nesting some group types within other group types, as shown in the following table. The table shows group scopes and limitations for nesting groups that are in the same domain.

Group scope	**Can nest Universal?**	**Can nest Global?**	**Can nest domain local?**
Domain local	Yes	Yes	Yes
Global	No	Yes	No
Universal	Yes	Yes	No

Convert groups including security, distribution, universal, domain local, and domain global

For the exam, you should study the following table which summarizes the limitations of group conversions.

Group scope	**Convert to Universal?**	**Convert to Global?**	**Convert to domain local?**
Domain local	Yes*	No	N/A
Global	Yes**	N/A	No
Universal	N/A	Yes***	Yes

*As long as other domain local groups are not members of the group

**As long as the group is not a member of another global group

***As long as other universal groups are not members of the group

Memorizing tables isn't ideal because memorizing all of the information is difficult. I recommend that you create a few groups in your lab and walk through the conversion process. It is helpful to click through and see conversions happening or see conversion failures.

Manage group membership using Group Policy

There are two ways to manage group membership by using Group Policy:

- **Restricted groups**. Restricted groups can control who is a member of a group and which groups a group is a member of.
- **Group Policy preferences - Local Users and Groups**. A second way to control group membership by using Group Policy is to use Group Policy preferences.

The following screen capture shows the restricted groups configuration windows in Group Policy. In the screen capture, a group named Group1 is being configured.

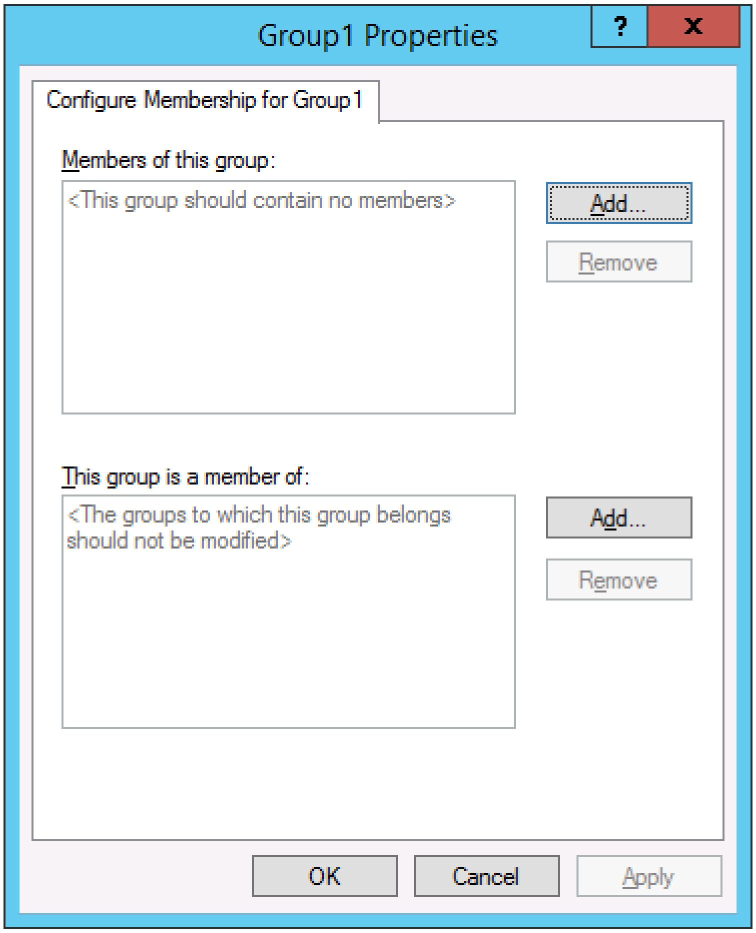

Group1 Properties

Configure Membership for Group1

Members of this group:

<This group should contain no members>

This group is a member of:

<The groups to which this group belongs should not be modified>

The top section controls who is a member of Group1. The key point to know is that whichever users and groups you specify in the Members of this group section will be the only members of the group. For example, if Group1 currently has Sue and Betty as members and you add Wanda and Wilma to the Members of this group section, only Wanda and Wilma will be members of Group1 after the GPO is processed.

The bottom section controls which group the specified group will be a member of. In this example, if you add Administrators to the This group is a member of section, then Group1 will be a member of the local Administrators group on all computers where the GPO applies.

You can use the top section to control membership of a group and the bottom section to control group nesting in a single GPO.

When you use Group Policy preferences, the configuration is a little different. The screen capture below shows a configuration where a domain group named IT is added to the local Administrators group. In this case, existing members of the local Administrators group remain as is.

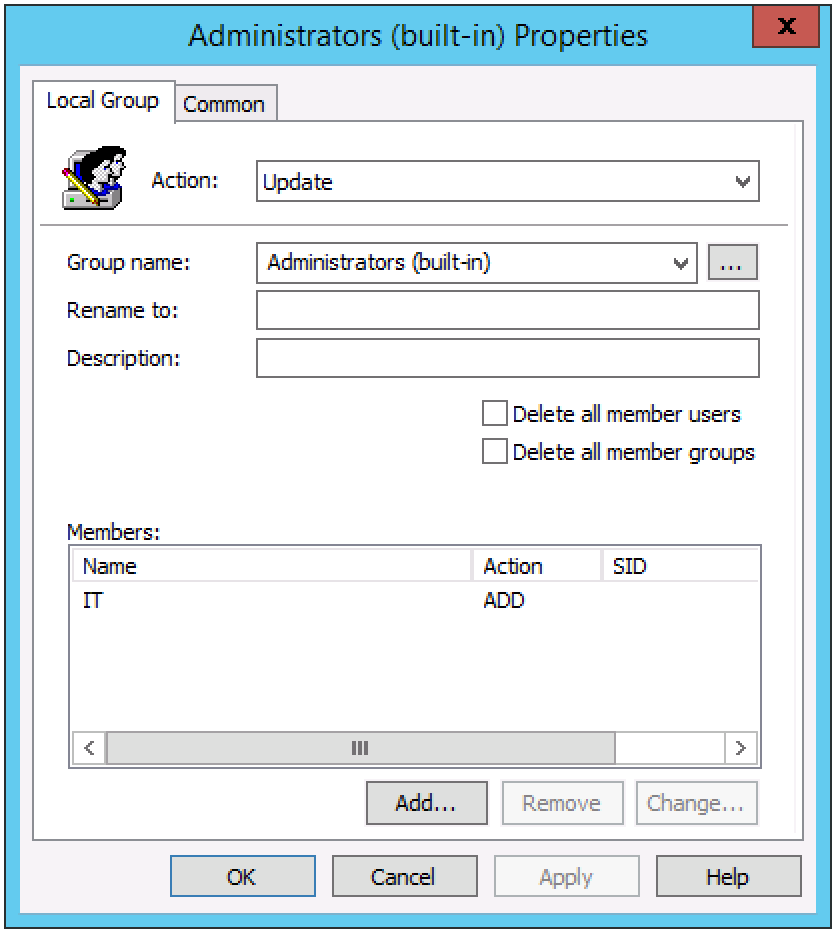

Administrators (built-in) Properties

Action: Update

Group name: Administrators (built-in)

Rename to:

Description:

□ Delete all member users

□ Delete all member groups

Members:

Name	Action	SID
IT	ADD	

You have the option to create, replace, update, or delete local groups. There is quite a bit of flexibility. On the exam, if you run across an item about controlling local group membership and the requirements are complex, Group Policy preferences is likely the right technology.

Enumerate group membership

There are several ways to look at the group membership of a user account, as described in the following examples. You should be familiar with all of these methods for the exam.

- **Use the whoami.exe command**. You can view the group memberships of the logged on user by running the **whoami /groups** command. If you are a member of a large number of groups, the output isn't ideal. But as a quick fix, especially on a computer where you don't have access to PowerShell, the whoami command is useful.

- **Use PowerShell to query group membership of a user**. You can view the group membership of a user named User1 in the Users container by running the **Get-ADUser "cn=User1,cn=Users,dc=adatum,dc=com" -Properties MemberOf | select -ExpandProperty MemberOf** command. The output is the DN of the groups.
- **Use PowerShell to find members of a group recursively**. You can view the group membership of a group named Domain Admins, including nested group membership, by running **the Get-ADGroupMember -Identity "Domain Admins" -Recursive | select Name** command. By removing the **-Recursive** parameter, you can view the direct membership of the group.

Delegate the creation and management of Active Directory objects

When delegating AD DS management tasks, you should opt to use the Delegation of Control Wizard when possible. It handles the majority of delegation tasks and makes it easy to recreate delegation, document the delegation process, and ensure consistency across environments.

The screen capture below shows the Delegation of Control Wizard delegating the following tasks:

- Create, delete, and manage user accounts.
- Create, delete, and manage groups
- Modify the membership of a group

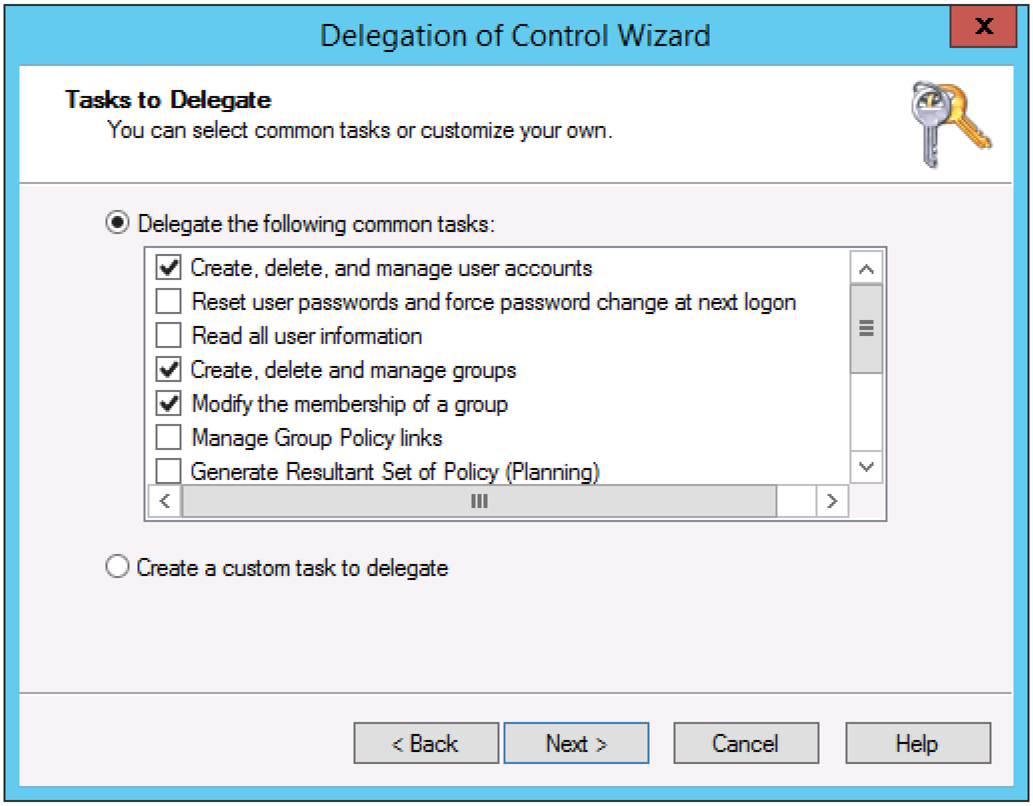

Tasks to Delegate
You can select common tasks or customize your own.

● Delegate the following common tasks:

☑ Create, delete, and manage user accounts
□ Reset user passwords and force password change at next logon
□ Read all user information
☑ Create, delete and manage groups
☑ Modify the membership of a group
□ Manage Group Policy links
□ Generate Resultant Set of Policy (Planning)

○ Create a custom task to delegate

Occasionally, you will need to delegate a custom task. When you need to delegate the rights for a user to update a single attribute or a few attributes, a custom task is probably your best choice to avoid giving users more rights than they need.

Know the following information for the exam:

- **The Delegation of Control Wizard is used to add permissions but cannot be used to remove permissions**. You can remove permissions using command-line tools or by manually modifying the ACLs.
- **You can use the dsacls command to view and modify permissions**. For the exam, you probably won't get tested on obscure dsacls command line parameters. However, you should be familiar with the basic command options. For example, /G is used to grant permissions while /R is used to delete permissions.
- **You can customize the Delegation of Control Wizard so that you can delegate permissions that aren't available by default**. You don't need to know more than that for the exam but if you are interested, take a look at Recipe 14.6 in the Active Directory Cookbook, 4th Edition.

Manage default Active Directory containers

The primary default containers that you manage in Active Directory are the Users container and the Computers container. Note that the Domain Controllers OU is not a container, although there was a time when it was referred to as a container. For the exam, now the difference between a container and an OU:

- **An OU can be targeted by Group Policy and can have child OUs**. For most organizations, all user objects and computer objects should be stored in OUs, not containers.
- **A container cannot be targeted by Group Policy and cannot have OUs under it**. By default, new user and computer objects are stored in the Users and Computers containers respectively, you can change the default location to use when new user and computer objects are created.

There are other default containers too. For the exam, the other container that you should be familiar with is the System container. By default, the System container is not visible in Active Directory Users and Computers but you can use the Advanced Features view to make it visible. Some of the backend objects such as password settings objects (used for fine-grained password policies) and the AdminSDHolder container (used to secure the ACLs on high security user objects) are stored in the System container.

Create, copy, configure, and delete groups and OUs

This section covers many of the routine tasks that administrators perform regularly. You should already be familiar with performing these tasks, at least by using the GUI tools. Here are some important things to know for the exam:

- **Be familiar with protecting objects from accidental deletion**. This option could be part of an exam scenario such as a troubleshooting scenario where you are not able to delete objects. By default, when you create a new OU, the object will be protected from accidental deletion. This is true whether you create the OU from the GUI or from PowerShell. However, other objects are not protected by default so you must use the option to protect them if needed.
- **Know what a subtree deletion is and when you have to use it**. If you try to delete an OU that contains other objects, you must use a subtree deletion or you cannot delete the OU. If you use a subtree deletion and the objects in the OU are protected from accidental deletion, they will still be deleted! When deleting with PowerShell, the -Recursive option will allow you to delete an OU with objects in it.
- **Know how to create groups and OUs**. The following PowerShell examples show the commands in use:
    - Create a new group named Group1: **New-ADGroup -Name Group1 -GroupScope Global**
    - Create a new OU named OU1 under the OU named POC in the adatum.com domain: **New-ADOrganizationalUnit -Name OU1 -Path “ou=POC,dc=adatum,dc=com”**

Some of the other tasks for this section such as managing group membership were discussed in previous sections.

6 - Create and manage Group Policy (16%)

Group Policy has a dedicated functional group on the 70-410 exam. As such, you should expect that it will be thoroughly covered on the exam. For the 70-410, you should expect operational questions with answers that have you performing tasks related to creating and configuring GPOs. If you do not work with Group Policy on a regular basis with your job, you should plan to perform all of the tasks in this section in a lab environment. Doing so will make a big difference in your retention of the information.

Create Group Policy Objects (GPOs)

Configure a Central Store

In Windows XP and Windows Server 2003, administrative template files (.adm files) are used for registry-based policy settings. You can import .adm files and then configure policy settings in a GPO. There are 2 downsides:

- **.ADM files are embedded in GPOs**. This takes up space and adds to the amount of data that has to be replicated during SYSVOL replication. The .ADM files are approximately 4MB in size. Imagine an environment that has 150 GPOs and multiple .ADM files.
- **When use Group Policy management tools, such as Group Policy Management Console (GPMC), all servers and client computers do not show the same administrative template settings**. Instead, you have to import .ADM files on all computers in order to see the settings. For example, imagine that you imported the Office 2010 administrative template file on Computer1 and used it to create a GPO to set some Office settings. A few days later, another administrator wants to edit the GPO to adjust some settings. He opens GPMC on Computer2 and edits the GPO. When he goes to find the Office 2010 settings, he doesn't see them because he doesn't have the administrative template file imported on Computer2. He can import it and proceed but it isn't a good administrative experience, especially when you have many administrators and many administrative template files.

A Central Store is a folder that you manually create in the SYSVOL folder on a domain controller. The folder name must be PolicyDefinitions. The following steps outline the process of creating a Central Store in a domain named contoso.com with a domain controller named DC01:

1. **Create the folder**. Create the PolicyDefinitions folder in \\dc01\sysvol\contoso.com\policies\ - thus, the path to the PolicyDefinitions folder is \\dc01\sysvol\contoso.com\policies\PolicyDefinitions. You

do not need to set permissions or create it more than one time. Once created, it will replicate to all of the other DCs during SYSVOL replication.

2. **Copy the needed files**. Copy all of the files from the PolicyDefinitions folder on a Windows server or client computer to the Central Store (\\dc01\sysvol\contoso.com\policies\PolicyDefinitions). There are .admx files (which are the .XML-based new version of .adm files) and .adml files (which are language-specific files so that settings are displayed in your desired language). For the exam, watch out for solutions that offer to only copy all of the .admx files to the Central Store. Also note that the language files are stored in a language specific folder. For US English, that folder is named en-US.

Once you have created a Central Store, you can go to the GPMC, edit a GPO, expand Policies (under Computer Configuration or User Configuration), and then expand Administrative Templates to verify that the Central Store is being used, as shown in the screen capture below:

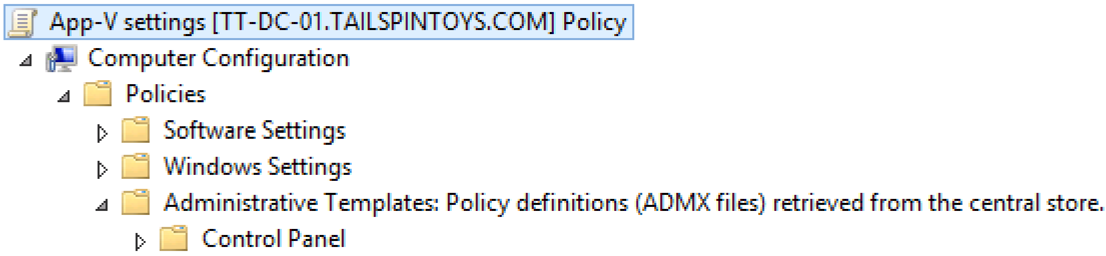

App-V settings [TT-DC-01.TAILSPINTOYS.COM] Policy
Computer Configuration
Policies
Software Settings
Windows Settings
Administrative Templates: Policy definitions (ADMX files) retrieved from the central store.
Control Panel

If the Central Store was not being used, you would see a reference to the local computer, as shown in the screen capture below:

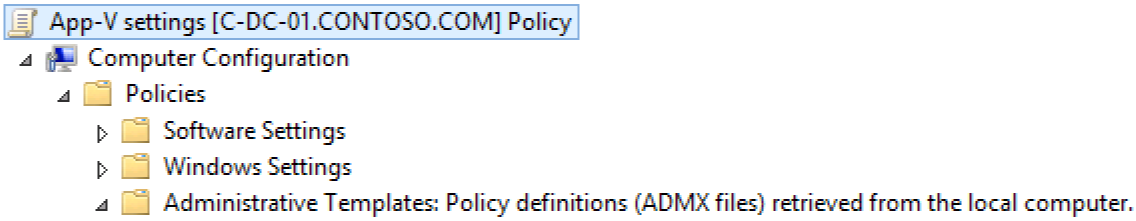

App-V settings [C-DC-01.CONTOSO.COM] Policy
Computer Configuration
Policies
Software Settings
Windows Settings
Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer.

There are a few things to keep in mind for the exam:

- **Watch for language specific questions**. For example, a scenario where a Central Store has been set up but the settings are not being displayed in the correct language. To fix this situation, you should copy the correct language specific .adml files to the Central Store.
- **Watch for legacy operating system questions**. Legacy operating systems (such as Windows XP or Windows Server 2003) cannot use the new .admx files. To fix this situation, you can upgrade the operating systems. Or, you can get rid of the Central Store and rely on the local PolicyDefinitions folder.

- **Watch for troubleshooting questions**. Administrators need to be able to access the files in the PolicyDefinitions folder. By default, the Authenticated Users group has Read & execute, List folder contents, and Read permissions to the Central Store's PolicyDefinitions folder. Also, be mindful of replication in troubleshooting questions. For example, you create a Central Store. An administrator in another site reports that when he opens the GPMC and edits a GPO, he is getting policy definitions from the local computer. To fix this problem, you should force replication and then have the administrator try again.

For more information, see the Managing Group Policy ADMX Files Step-by-Step Guide.

Manage starter GPOs

Starter GPOs are GPO templates that you can use when you create a new GPO. There are ten default starter GPOs. You can edit those or create your own. When you create a new GPO from a Starter GPO, the new GPO will be configured with the same settings as the Starter GPO. Here are some examples of use cases for Starter GPOs:

- **You have 15 sites**. You plan to use the same client computer firewall settings across all client computers. You plan to use site-specific settings for backgrounds, IE settings, and other miscellaneous settings. In this case, you could create a Starter GPO that has all of the client computer firewall settings and use that to create the 15 site-specific GPOs.
- **You have multiple domains or forests**. You plan to use the same settings in multiple domains and want to ensure consistency across the GPOs. You can export starter GPOs and import them into other domains. The export process saves a starter GPO as a cabinet (.cab) file.

By default, starter GPOs aren't displayed in the GPMC. Instead, you see a message indicating that you need to create the folder, as shown in the screen capture below.

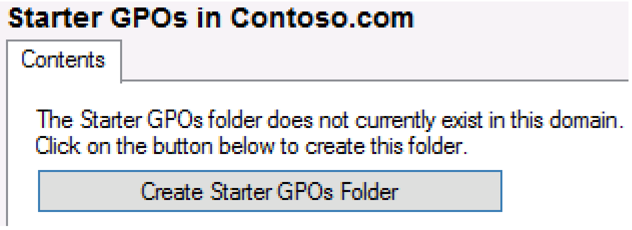
Starter GPOs in Contoso.com

The Starter GPOs folder does not currently exist in this domain.
Click on the button below to create this folder.

After you create the Starter GPOs folder, you see the ten default starter GPOs, as shown in the screen capture below.

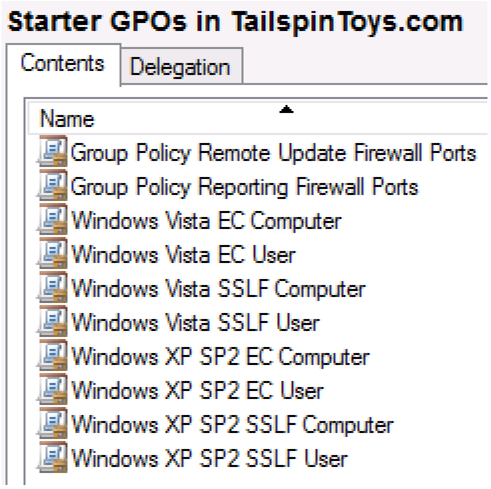

Starter GPOs in TailspinToys.com

Group Policy Remote Update Firewall Ports
Group Policy Reporting Firewall Ports
Windows Vista EC Computer
Windows Vista EC User
Windows Vista SSLF Computer
Windows Vista SSLF User
Windows XP SP2 EC Computer
Windows XP SP2 EC User
Windows XP SP2 SSLF Computer
Windows XP SP2 SSLF User

You can click the Create Starter GPOs Folder by clicking the button. To create a starter GPO named GPO1 by using PowerShell, you need to run the **New-GPStarterGPO -Name GPO1** command. The only other related PowerShell cmdlet for starter GPOs is Get-GPStarterGPO.

For the exam, here are a few key points about starter GPOs.

- You only need to create the Starter GPOs folder on one DC.
- Both of the starter GPO cmdlets are in the module named GroupPolicy. If you see a troubleshooting question related to starter GPOs and PowerShell, remember that the GroupPolicy module may need to be imported.
- To export a starter GPO, you click **Save as Cabinet** while highlighting the starter GPO in GPMC.
- To import a starter GPO, you click **Load Cabinet** while highlighting the Starter GPOs container in GPMC.

Configure GPO links

The primary tasks you perform when configuring GPO links is creating links and removing links. While most administrators are familiar with clicking a GPO and dragging it to an OU to link it to the OU, you should be familiar with the less-known methods too:

- **PowerShell**. To link a GPO named GPO9 to the domain level in a domain named contoso.com, you can run the **New-GPLink -Name GPO9 -Target “dc=contoso,dc=com”** command. You can use the **-LinkEnabled** and **-Order** parameters to configure the link too. Remember, the order of the link establishes when the settings are applied. The last GPO that is applied takes precedence. A GPO with a link order of one is the highest precedence GPO and will be applied last.

- **Create a new GPO and link it at the time of creation**. You can right-click an OU and then choose the option to create a new GPO and link it at that OU.

In addition to knowing how to create and delete links, you should clearly understand some of the characteristics of GPO links:

- **Memorize LSDOU**. LSDOU is an acronym that follows the GPO processing order. Group Policy is applied at the **L**ocal computer level first (local GPO), at the **S**ite level second, at the **D**omain level third, and at the **OU** level last.
- **GPOs linked to sites aren't easy to notice**. Watch for troubleshooting scenarios where you are having unexpected GPO application and the GPO isn't linked to any OUs. The GPO may be linked to a site.
- **GPOs linked to parent OUs are processed before GPOs linked to their child OUs**. This small tidbit of info is handy to know for troubleshooting situations.

Configure multiple local group policies

Multiple Local Group Policy objects (MLGPO) was introduced with Windows Vista. Prior to Windows Vista, administrators could use a single GPO for a standalone computer (a computer not joined to a domain). MLGPO allows administrators to apply the following local policies:

- **Local Group Policy**. This is the only policy that you can use to configure computer settings (although you can also configure user settings). This policy is applied first. Thus, user settings can be overwritten by other local policies.
- **Local Computer\Administrators policy**. This policy is applied to users in the local Administrators group. You can only specify user settings in this policy. This policy applies second, after the Local Group Policy object. If you need to ensure that user settings are going to be applied to administrative users, you should configure the settings in this policy.
- **Local Computer\Non-Administrators policy**. This policy is applied to all local users that are not members of the local Administrators group. You can only specify user settings in this policy. This policy applies last, after all other local policies. If you need to ensure that user settings are going to be applied to non-administrative users, you should configure the settings in this policy.
- **Local Computer\<user> policy**. In this case, the username of the user would be shown instead of "<user>". This policy only applies to a single user. Thus, you could have several policies with each one only applying to a specific user. If you need settings to apply to all local users, you should use the Local Computer\Non-Administrators policy because it applies to multiple users so it is a more efficient way of managing settings.

Standalone computers, such as kiosks and often DMZ servers, can benefit from local policies. Watch for exam situations involving high security environments or environments where computers aren't often joined to a domain (training centers, kiosks, small business, DMZ servers).

Configure security filtering

By default, all GPOs have a security filter set to the Authenticated Users group. To ensure that a GPO only applies to a smaller subset of objects, you can change the security filtering on the GPO.

Below is a screen capture of the default security filtering.

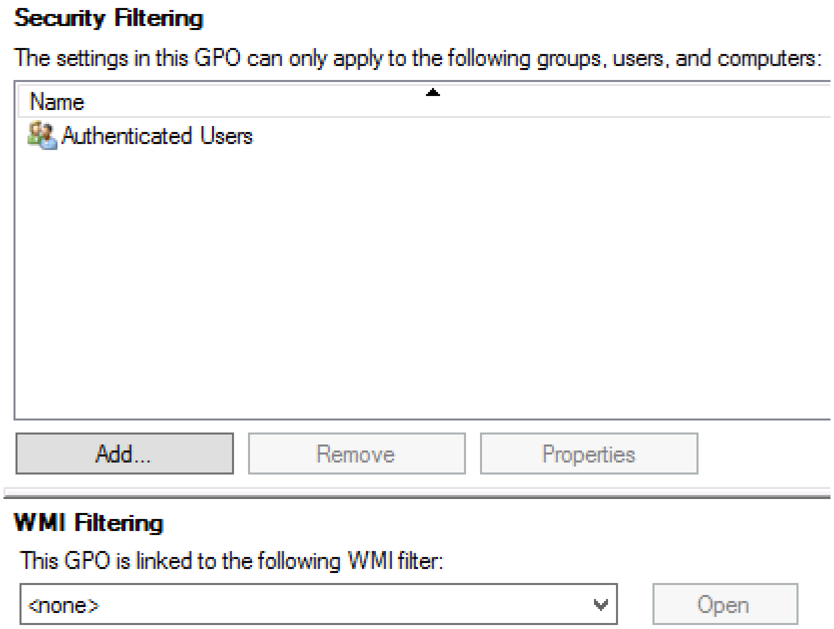

Security Filtering

The settings in this GPO can only apply to the following groups, users, and computers:

Name
Authenticated Users

WMI Filtering

This GPO is linked to the following WMI filter: <none>

When you click the Add button, you cannot specify the permissions. Instead, the user or group that you add is given Read and Apply Group Policy permissions. Alternatively, you can have more granular control by using the Delegation tab.

Familiarize yourself with the PowerShell commands for working with security filtering. The following are some examples:

- Give a User1 Read and Apply Group Policy permissions on a GPO named GPO1: **Set-GPPermission -Name "GPO1" -TargetName "User1" -TargetType User -PermissionLevel GpoApply**

- Give the GPO_Admins group the ability to edit all GPOs in the domain: **Set-GPPermission -All -TargetName "GPO_Admins" -TargetType Group -PermissionLevel GpoEdit -Replace** (be sure to remember the **-All** parameter and know that the **-Replace** parameter replaces existing permissions)
- Find out which permissions the Authenticated Users group has on GPO1: **Get-GPPermission -Name "GPO1" -TargetName "Authenticated Users" -TargetType Group | FL Permission**

Here are the key things to remember for the exam regarding security filtering:

- **In order for a user or a computer to process a GPO, they need Read and Apply Group Policy permissions**. Watch for exam scenarios where you are troubleshooting GPO issues and you add a user or group and give them Read permissions to the GPO. You also need to give the Apply Group Policy permission.
- **The Authenticated Users group contains authenticated users and authenticated computers**. This is an important fact. Watch for an exam scenario where you are troubleshooting GPO issues to a computer and the computer isn't processing the policy. An answer that has you give the computer object Read and Apply Group Policy permissions would be very compelling if you didn't know that the Authenticated Users group contains users and computers.
- **Know the PowerShell subtleties**. The GpoApply permission level in PowerShell gives Read permission too. Also, there are two aliases: Get-GPPermissions and Set-GPPermissions. These are aliases to Get-GPPermission and Set-GPPermission.

Configure security policies

Every Windows server has a local security policy that you can edit by clicking Local Security Policy in the Tools menu of Server Manager. The local security policy is a portion of the local Group Policy that handles security-related settings. You can modify the same settings in the local policy by navigating to Computer Configuration/Windows Settings/Security Settings. All of the Group Policy processing information about local policy, discussed previously, applies to the local security policy since it is just component of local policy.

Configure User Rights Assignment

The User Rights Assignment section of a GPO has 44 settings that you can use to configure user-related security settings:

Access Credential Manager as a trusted caller	Access this computer from the network	Act as part of the operating system
Add workstations to domain	Adjust memory quotas for a process	Allow log on locally
Allow log on through Remote Desktop Services	Back up files and directories	Bypass traverse checking
Change the system time	Change the time zone	Create a pagefile
Create a token object	Create global objects	Create permanent shared objects
Create symbolic links	Debug programs	Deny access to this computer from the network
Deny log on as a batch job	Deny log on as a service	Deny log on locally
Deny log on through Remote Desktop Services	Enable computer and user accounts to be trusted for delegation	Force shutdown from a remote system
Generate security audits	Impersonate a client after authentication	Increase a process working set
Increase scheduling priority	Load and unload device drivers	Lock pages in memory
Log on as a batch job	Log on as a service	Manage auditing and security log
Modify an object label	Modify firmware environment values	Perform volume maintenance tasks
Profile single process	Profile system performance	Remove computer from docking station
Replace a process level token	Restore files and directories	Shut down the system
Synchronize directory service data	Take ownership of files or other objects	

From an exam perspective, it is likely that this section will cover using one or more of the settings. It isn't likely that you'll run into an archaic or rarely used setting. But, spend a little time looking through the settings on a live system so that you can see which settings are configured by default and which aren't. If you have extra time, you can open each setting in a GPO and read the information in the Explain tab. The settings that deny access can easily be used to create troubleshooting questions for the exam. None of those are configured by default. A few other important settings to be familiar with are:

- **Log on as a service**. If you want to use a service account to run a Windows service, it must be given the right to log on as a service. An exam item could present a situation where you have a service account, it was added to the local Administrators group, you configure it for a service, but the service won't start. Look for an answer that gives the service account rights to log on as a service.
- **Allow log on locally**. If a user or administrator needs to log on locally, at the console, to a computer, then they need to have the ability to log on locally. Similar to other settings, the local Administrators group has this by default. Thus, watch for answer choices that are redundant and have you adding a user to this setting even though they are already in the local Administrators group (since such an action won't change anything). Also beware of exam scenarios asking you to maximize security as part of the solution. In such cases, be careful – sometimes the answer could still be giving out more permissions or rights than required (because first and foremost, the answer has to meet the requirement).
- **Allow log on through Remote Desktop Services**. This group gives administrators the right to connect to a computer by using Remote Desktop Connection over RDP. By default, the Administrators group and the Remote Desktop Users group have this right. In a scenario where a user cannot connect to a computer over RDP, you could fix the problem by adding the user to Administrators, Remote Desktop Users, or by directly giving the user the **Allow log on through Remote Desktop Services** right.

Configure Security Options settings

The Security Options section is the final section of the dedicated security settings area of a GPO. The following settings are available:

Accounts: Administrator account status	Accounts: Block Microsoft accounts	Accounts: Guest account status	Accounts: Limit local account use of blank passwords to console logon only
Accounts: Rename administrator account	Accounts: Rename guest account	Audit: Audit the access of global system objects	Audit: Audit the use of Backup and Restore privilege

Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings	Audit: Shut down system immediately if unable to log security audits	DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax	DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax
Devices: Allow undock without having to log on	Devices: Allowed to format and eject removable media	Devices: Prevent users from installing printer drivers	Devices: Restrict CD-ROM access to locally logged-on user only
Devices: Restrict floppy access to locally logged-on user only	Domain controller: Allow server operators to schedule tasks	Domain controller: LDAP server signing requirements	Domain controller: Refuse machine account password changes
Domain member: Digitally encrypt or sign secure channel data (always)	Domain member: Digitally encrypt secure channel data (when possible)	Domain member: Digitally sign secure channel data (when possible)	Domain member: Disable machine account password changes
Domain member: Maximum machine account password age	Domain member: Require strong (Windows 2000 or later) session key	Interactive logon: Display user information when the session is locked	Interactive logon: Do not display last user name
Interactive logon: Do not require CTRL+ALT+DEL	Interactive logon: Machine account lockout threshold	Interactive logon: Machine inactivity limit	Interactive logon: Message text for users attempting to log on
Interactive logon: Message title for users attempting to log on	Interactive logon: Number of previous logons to cache (in case domain controller is not available)	Interactive logon: Prompt user to change password before expiration	Interactive logon: Require Domain Controller authentication to unlock workstation
Interactive logon: Require smart card	Interactive logon: Smart card removal behavior	Microsoft network client: Digitally sign communications (always)	Microsoft network client: Digitally sign communications (if server agrees)
Microsoft network client: Send unencrypted password to third-party SMB servers	Microsoft network server: Amount of idle time required before suspending session	Microsoft network server: Attempt S4U2Self to obtain claim information	Microsoft network server: Digitally sign communications (always)
Microsoft network server: Digitally sign communications (if client agrees)	Microsoft network server: Disconnect clients when logon hours expire	Microsoft network server: Server SPN target name validation level	Network access: Allow anonymous SID/Name translation
Network access: Do not allow anonymous enumeration of SAM accounts	Network access: Do not allow anonymous enumeration of SAM accounts and shares	Network access: Do not allow storage of passwords and credentials for network authentication	Network access: Let Everyone permissions apply to anonymous users

Network access: Named Pipes that can be accessed anonymously	Network access: Remotely accessible registry paths	Network access: Remotely accessible registry paths and sub-paths	Network access: Restrict anonymous access to Named Pipes and Shares
Network access: Shares that can be accessed anonymously	Network access: Sharing and security model for local accounts	Network security: Allow Local System to use computer identity for NTLM	Network security: Allow LocalSystem NULL session fallback
Network security: Allow PKU2U authentication requests to this computer to use online identities.	Network security: Configure encryption types allowed for Kerberos	Network security: Do not store LAN Manager hash value on next password change	Network security: Force logoff when logon hours expire
Network security: LAN Manager authentication level	Network security: LDAP client signing requirements	Network security: Minimum session security for NTLM SSP based (including secure RPC) clients	Network security: Minimum session security for NTLM SSP based (including secure RPC) servers
Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication	Network security: Restrict NTLM: Add server exceptions in this domain	Network security: Restrict NTLM: Audit Incoming NTLM Traffic	Network security: Restrict NTLM: Audit NTLM authentication in this domain
Network security: Restrict NTLM: Incoming NTLM traffic	Network security: Restrict NTLM: NTLM authentication in this domain	Network security: Restrict NTLM: Outgoing NTLM traffic to remote servers	Recovery console: Allow automatic administrative logon
Recovery console: Allow floppy copy and access to all drives and all folders	Shutdown: Allow system to be shut down without having to log on	Shutdown: Clear virtual memory pagefile	System cryptography: Force strong key protection for user keys stored on the computer
System cryptography: Use FIPS compliant algorithms for encryption	System objects: Require case insensitivity for non-Windows subsystems	System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links)	System settings: Optional subsystems
System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies	User Account Control: Admin Approval Mode for the Built-in Administrator account	User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop	User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode
User Account Control: Behavior of the elevation prompt for standard users	User Account Control: Detect application installations and prompt for elevation	User Account Control: Only elevate executables that are signed and validated	User Account Control: Only elevate UIAccess applications that are installed in secure locations

User Account Control: Run all administrators in Admin Approval Mode	User Account Control: Switch to the secure desktop when prompting for elevation	User Account Control: Virtualize file and registry write failures to per-user locations	

There are a lot of security options. In this study guide, we will not dive into the settings because memorizing what they all do isn't a good use of study time since there is likely only to be a couple of questions on the exam about these settings. However, it is a good idea to look at the Explain tab for all of the settings. Give them a quick read. Then move on.

Configure Security templates

Microsoft offers quite a bit of security guidance for securing your environment. Part of the guidance comes in the form of security baselines. You can download the latest guidance and baselines here. While you can directly download and work with the provided baselines, it is a good idea to use Security Compliance Manager (SCM) instead. SCM is a free tool from Microsoft to make working with baselines and automating GPO creation a much easier task. For the exam, it is highly likely that SCM will be the focus of questions in this section.

For the exam, you should download and install SCM. You can download it here. You should know how to perform the following tasks in SCM:

- **Download new baselines**. By default, SCM does not have baselines for Windows 8.1, Windows Server 2012 R2, or other newer products. You must download and import them after installation.
- **Edit baselines**. You must have a modifiable copy of each baseline in order to make any changes to it. You can do this during the import or do this individually by duplicating baselines in SCM after import.
- **Use baselines to create GPOs**. You should be familiar with the process of taking a baseline, customizing it, and then implementing it into a GPO:
    - Duplicate the baseline.
    - Customize the baseline.
    - Export the baseline using the **GPO Backup (folder)** option.
    - Create a new GPO.
    - Import the settings from the exported baseline (GPO backup). Note that you cannot create a new GPO and then restore it from the GPO backup of the exported baseline. While that sounds like a really good exam answer, you can't do it!

Below is a screen capture of the import process. Note the **Create modifiable copies of each baseline to be imported** option because it is an important option that you need to use so that you can edit the baselines.

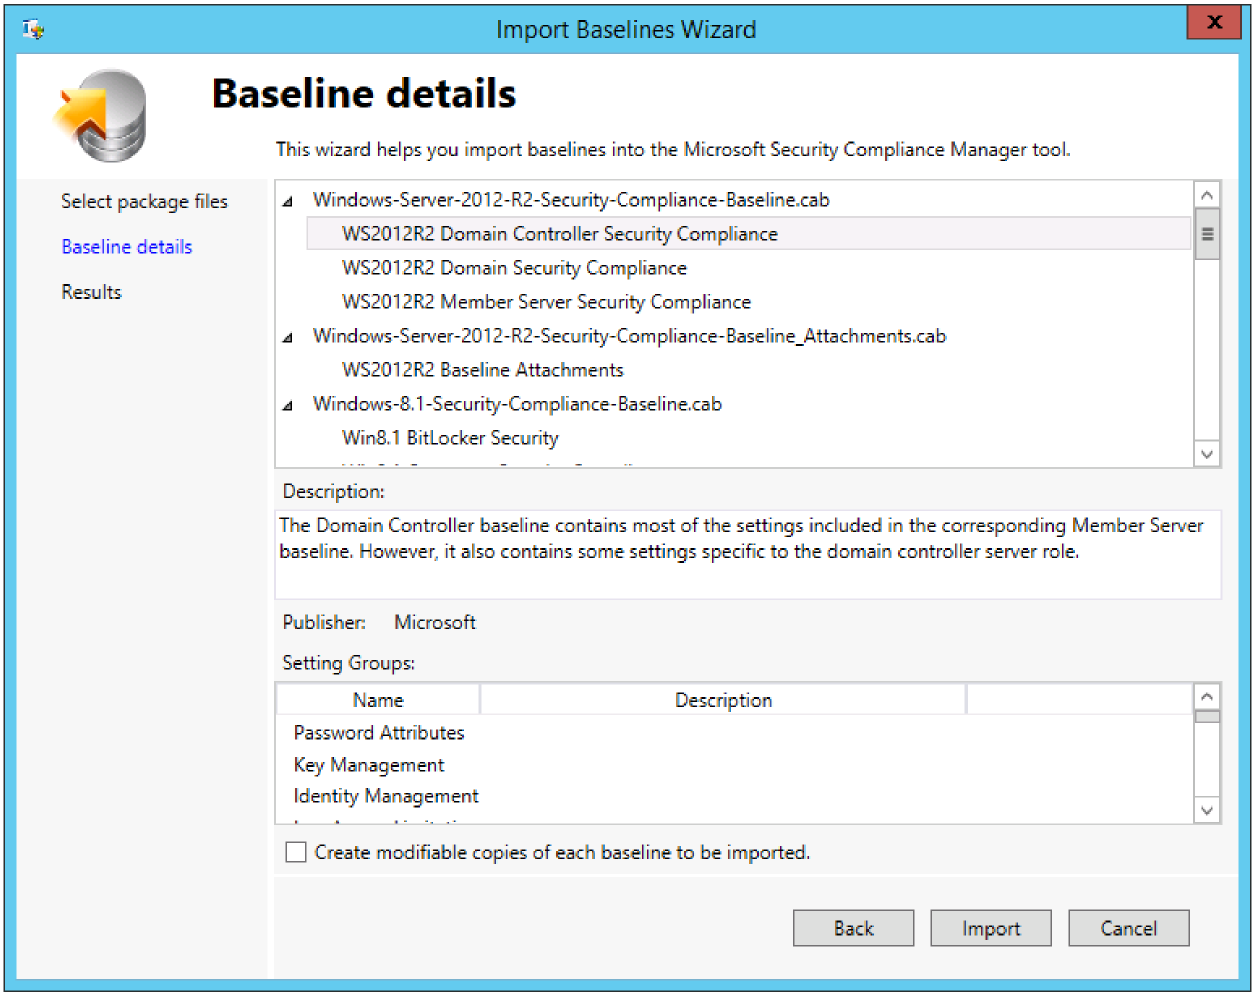

Baseline details

This wizard helps you import baselines into the Microsoft Security Compliance Manager tool.

- Windows-Server-2012-R2-Security-Compliance-Baseline.cab
  - WS2012R2 Domain Controller Security Compliance
  - WS2012R2 Domain Security Compliance
  - WS2012R2 Member Server Security Compliance
- Windows-Server-2012-R2-Security-Compliance-Baseline_Attachments.cab
  - WS2012R2 Baseline Attachments
- Windows-8.1-Security-Compliance-Baseline.cab
  - Win8.1 BitLocker Security

Description:

The Domain Controller baseline contains most of the settings included in the corresponding Member Server baseline. However, it also contains some settings specific to the domain controller server role.

Publisher: Microsoft

Setting Groups:

Name	Description
Password Attributes	
Key Management	
Identity Management	

☐ Create modifiable copies of each baseline to be imported.

Configure Audit Policy

Most administrators are familiar with the basic auditing settings which have been around for a long time:

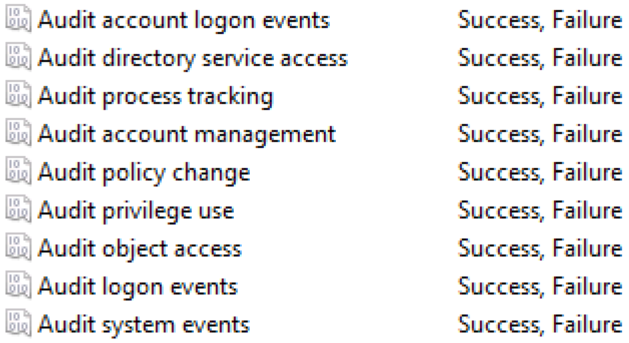

Audit account logon events	Success, Failure
Audit directory service access	Success, Failure
Audit process tracking	Success, Failure
Audit account management	Success, Failure
Audit policy change	Success, Failure
Audit privilege use	Success, Failure
Audit object access	Success, Failure
Audit logon events	Success, Failure
Audit system events	Success, Failure

For the exam, you should be familiar with these audit settings and understand which settings you need to enable to capture specific types of information.

Additionally, you need to understand the difference between these basic audit settings and the newer advanced audit categories:

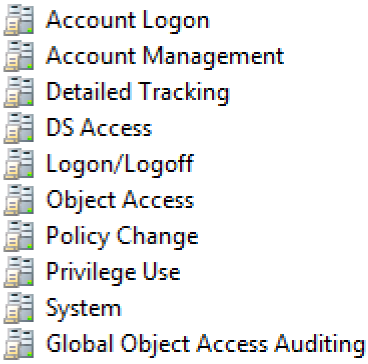

Account Logon
Account Management
Detailed Tracking
DS Access
Logon/Logoff
Object Access
Policy Change
Privilege Use
System
Global Object Access Auditing

You do not individually configure auditing at the category level. Instead, you can expand each category and set granular audit settings in the category. With advanced auditing, you can figure 58 auditing settings compared to nine with the basic auditing settings. Advanced auditing allows administrators to audit exactly what they need without having to capture audit information that they don't need. For example, you can configure auditing so that you capture account lockout events but no other account logon events. Be aware of the following characteristics of auditing:

- **Which one takes precedence**. If you configure auditing with the basic audit settings and you configure auditing with the advanced audit policy settings, the advanced audit policy settings take precedence.
- **Global Object Access Auditing is a newer feature that allows you to configure auditing on file system or registry objects globally**. Prior to global object access auditing, if you wanted to audit a folder, you had to enable object access auditing and then go change the System Access Control List (SACL) on the folder to turn the auditing on. Now, with global object access auditing, you can take care of the SACL work and use the object access auditing to enable auditing. With this method, you don't have to configure anything for the computer that contains the files and folders to be audited.
- **The more you audit, the more information you capture**. And that leads to using up a large amount of disk space along with potential performance problems. Watch for exam scenarios where you need to maximize performance with auditing. The answer should be using the advanced audit policy settings and granularly selecting only what you need to audit.

Configure Local Users and Groups

This topic is a little bit ambiguous. Not just for test takers. But also for exam item writers. They have to interpret the OD just like test takers. The functional group is titled “Create and manage Group Policy” so that indicates that everything under it will be related to Group Policy. The topic is under the “Configure security policies” objective which, at first glance, makes it seem like it is related to the local security policy. However, it isn’t. Because there aren’t any settings related to configuring local users and groups in the local security policy. For this topic, you should know the following information:

- **How to work with Restricted Groups**. Restricted groups, in the context of this exam topic, help you automatically add users or groups into local groups on a group of computers. For example, if you want to add the Server_Admins group in your domain to the local Administrators group on all of your servers, you could use Restricted Groups. The screen capture below shows the configuration window for doing this. Be aware of the two options:
  - Members of this group. This option controls which users or computers are members of the specified group.
  - This group is a member of. This option control which other groups the specified group is a member of. This is the option used in the screen capture below.

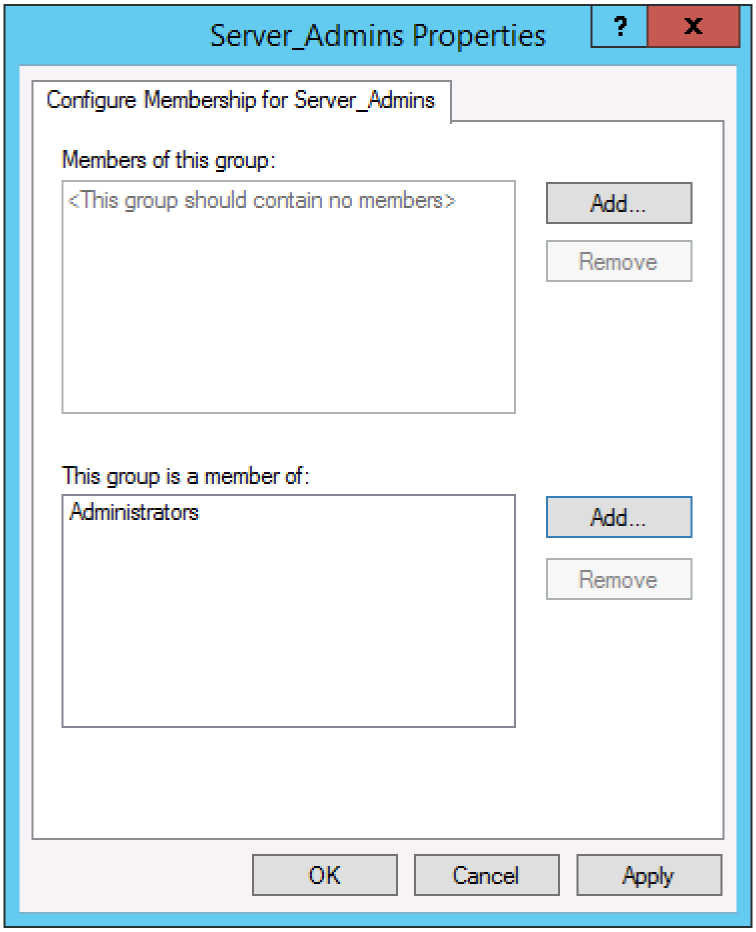

Members of this group:

<This group should contain no members>

This group is a member of:

Administrators

You can use enforce who is a member of a group and which groups the group belongs to in a single policy.

- **Restricted groups works for domain users and groups as well as local users and groups**. If you configure a policy to add a group named Server_Admins to a group named Administrators and you link the policy to the domain, the Server_Admins group will be a member of the Administrators group in the domain as well as the local Administrators group on member computers. Thus, it is important to link the policy to the correct location.
- **You can also use Group Policy Preferences to automatically add domain-based users or groups to local groups on member computers**. The settings are located under Computer Configuration, Preferences, Control Panel Settings, Local Users and Groups. The configuration to add Server_Admins to the local Administrators group on member computers is shown in the screen capture below.

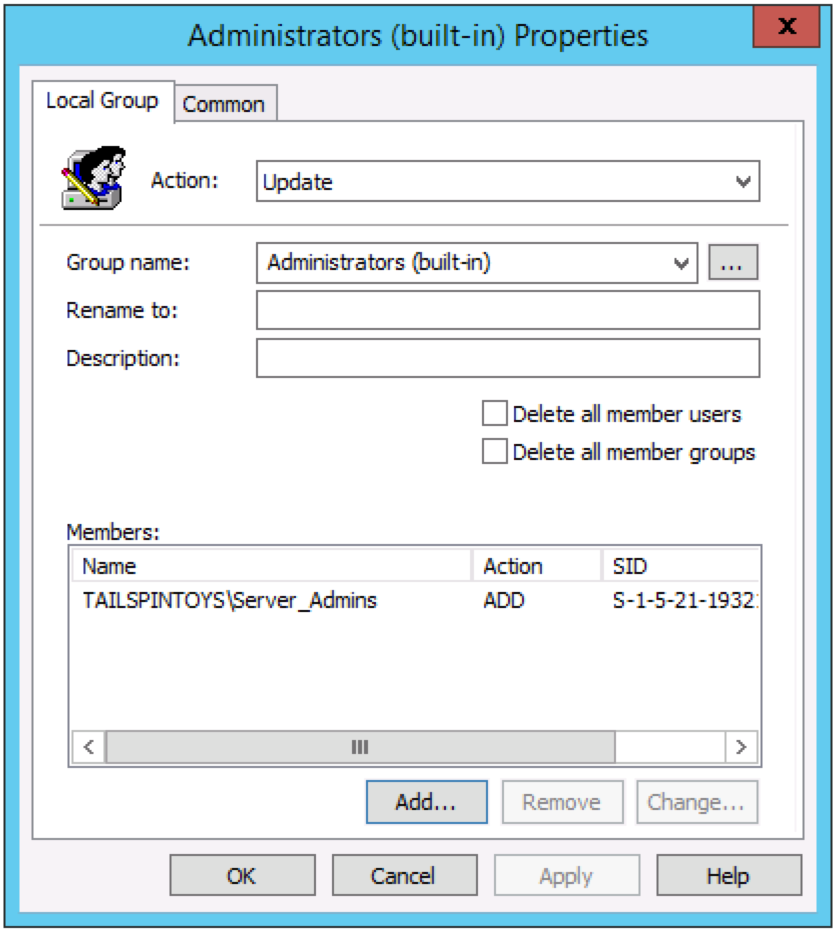

Action: Update

Group name: Administrators (built-in)

Rename to:

Description:

☐ Delete all member users

☐ Delete all member groups

Members:

Name	Action	SID
TAILSPINTOYS\Server_Admins	ADD	S-1-5-21-1932

- You can create, update, replace, or delete local User accounts by using Group Policy.
- You can create, update, replace, or delete local groups by using Group Policy.

Configure User Account Control (UAC)

You can configure UAC by using Group Policy. The settings are located in the Computer Configuration/Security Settings/Security Options node. There are 10 settings available and you should be familiar with the settings:

UAC GPO setting name	Description
User Account Control: Admin Approval Mode for the Built-in Administrator account	This setting dictates whether the local Administrator account will be subject to UAC prompts or will not be subject to UAC prompts.
User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop	This setting, if enabled, allows some applications (such as remote control applications) to allow the UAC prompt to be displayed to the remote user instead of just to the local user on the secure desktop.
User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode	This setting dictates whether administrators are prompted for UAC and how such prompts are shown (ask for consent, ask for credentials, and more).
User Account Control: Behavior of the elevation prompt for standard users	This is the same as the UAC setting for administrators directly above but used for standard users, not administrators.
User Account Control: Detect application installations and prompt for elevation	If enabled, you will be prompted during an installation if it requires privilege escalation.
User Account Control: Only elevate executables that are signed and validated	If enabled, PKI signature checks are performed when elevation is required.
User Account Control: Only elevate UIAccess applications that are installed in secure locations	If enabled, applications that require the User Interface Accessibility (UIAccess) integrity level must be stored in a secure location such as Program Files, Windows\System32, and Program Files (x86).
User Account Control: Run all administrators in Admin Approval Mode	This setting, if enabled, enforces Admin Approval Mode where all administrative users are set to use UAC.
User Account Control: Switch to the secure desktop when prompting for elevation	This setting controls whether the secure desktop is used or the interactive desktop is used for UAC prompts. For high security environments, you should enable this setting.
User Account Control: Virtualize file and registry write failures to per-user locations	If you disable this setting, applications that write data to protected locations fail. If you enable this setting, write failures are redirected to defined user locations.

Configure application restriction policies

When you think of application restriction policies, think of Software Restriction Policies (SRP) and AppLocker. The exam will mostly focus on AppLocker because it is the newest version of Microsoft's application restriction technology. AppLocker was introduced with Windows 7 and Windows Server 2008 R2. It is important to note the differences between the two technologies because it is likely that you will see some reference to Software Restriction Policies in questions with legacy environments or in use as compelling answer choices. Know the following information for the exam:

- **If you use SRP and AppLocker to target a computer, SRP is ignored**. What this means is that AppLocker takes precedence when both technologies are targeting the same computer. However, you can use SRP for legacy clients such as Windows Server 2003 and Windows XP while simultaneously using AppLocker for newer clients. It is a good practice to use a separate GPO for each technology.
- **AppLocker can control packaged apps and packaged app installers**. AppLocker began to support control over packaged apps with Windows Server 2012 and Windows 8. SRP does not support packaged apps.
- **AppLocker supports rule exceptions**. You can create a rule to Allow everything in the Program Files directory to run except files by a specific publisher, in a specific subfolder, or with a specific file hash.

You can read additional detail comparing SRP and AppLocker here.

Configure rule enforcement

There are 2 primary enforcement options in AppLocker:

- **Enforced**. When rule enforcement is configured, rules are enforced and cannot be overridden.
- **Audit only**. When you configure rule enforcement for auditing, rules are not enforced. This mode is helpful when first configuring AppLocker because it shows you the impact of your rules before you begin enforcing them. You can check a dedicated Windows Event Log under the AppLocker folder to learn how the rules are working in audit mode (there are four dedicated log files under there - one for EXE and DLL, one for MSI and Script, one for Packaged app-Deployment, and one for Packaged app-Execution). From the exam perspective, watch for scenarios that want to minimize risk or not impact functionality because audit mode could be the way to go in such scenarios.

Configure AppLocker rules

An important concept when working with rules is knowing about the automatic rules that AppLocker will create for you. When you configure AppLocker in a GPO, you can right-click and then automatically generate rules, as shown in the screen capture below.

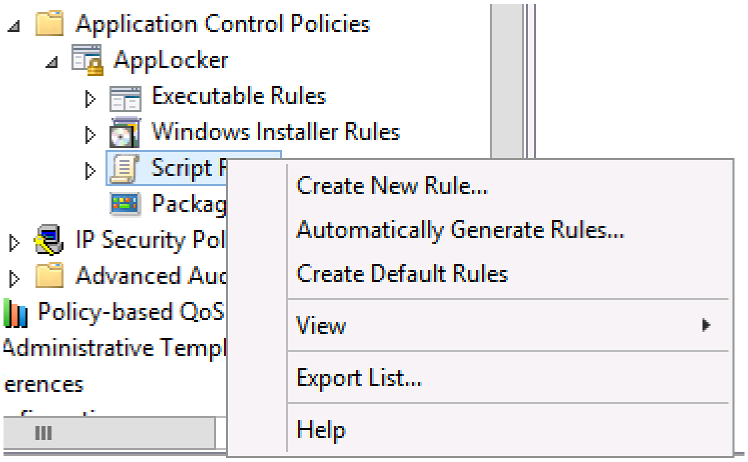

Automatically generating rules will scan a specified folder and create rules based on your input.

You can also create the default rules. Default rules are available for each of the rule types: Executable, Windows Installer, Script, and Packaged app. You shouldn't plan on memorizing all of the default rules or the automatic rules that would be created in a default installation of Windows. However, you should know the following about the default rules:

Executable:

- The default rules allow the Everyone group to run executable files located in the Program Files folder.
- The default rules allow the Everyone group to run executable files located in the Windows folder.
- The default rules allow the local Administrators group to run all executable files regardless of location.

Windows Installer:

- The default rules allow the Everyone group to run all digitally signed Windows Installer files.
- The default rules allow the Everyone group to run all Windows Installer files in the %systemdrive%\Windows\Installer folder.
- The default rules allow the local Administrators group to run all Windows installer files regardless of location.

Script:

- The default rules allow the Everyone group to run scripts located in the Program Files folder.
- The default rules allow the Everyone group to run scripts located in the Windows folder.
- The default rules allow the local Administrators group to run all scripts regardless of location.

Packaged app:

- The default rules allow the Everyone group to run all signed packaged apps.

From an exam perspective, watch for problems or goals that could be solved by creating the default rules or by automatically creating rules. This could be a troubleshooting scenario or a new deployment. For example, imagine a scenario when another administrator implements AppLocker and thereafter all users are able to run everything in the Program Files folder. Now, imagine that a company requirement for the Telemarketing department is that they should only be able to run a specific application. You could delete the default rules. Then, you could add a rule for a Telemarketing group and the specific application. Or, imagine a scenario where you set up AppLocker but then installed a bunch of applications after that. You could automatically create rules so that users can access the newly installed apps.

On the exam, watch for scenarios where you need to minimize administrative overhead because that may indicate an answer involving automatic rules or default rules.

Configure Software Restriction Policies

To prepare for the exam, you should take a look at SRP in GPMC. Edit a GPO and navigate to Computer Configuration/Policies/Windows Settings/Security Settings/Software Restriction Policies. Often, visualizing the available settings (and even better is using them) along with reading some key points will maximize your information retention. Know the following key points about SRPs:

- You can set SRP to Disallowed. By doing so, users cannot run software.
- You can set SRP to Basic User. By doing so, non-Administrative users can run programs.
- You can set SRP to Unrestricted. This is the default setting. Access to run programs is determined by the access rights to the program executable and supporting files.
- SRP allows you to you restrict access to programs or the registry.
- SRP is only valid for Windows XP and Windows Server 2003.

SRP enforcement allows you to configure SRP so it applies to all users or only non-administrative users. You can enforce or ignore certificates. You can apply SRP to all files or all files except library files such as .DLL files. Note that choosing to use SRP on all files could create some compatibility issues so watch for scenarios where there are issues after implementing SRP because in such a situation, configuring SRP to work with all files except library files may be a fix. The screen capture below shows the default SRP enforcement settings:

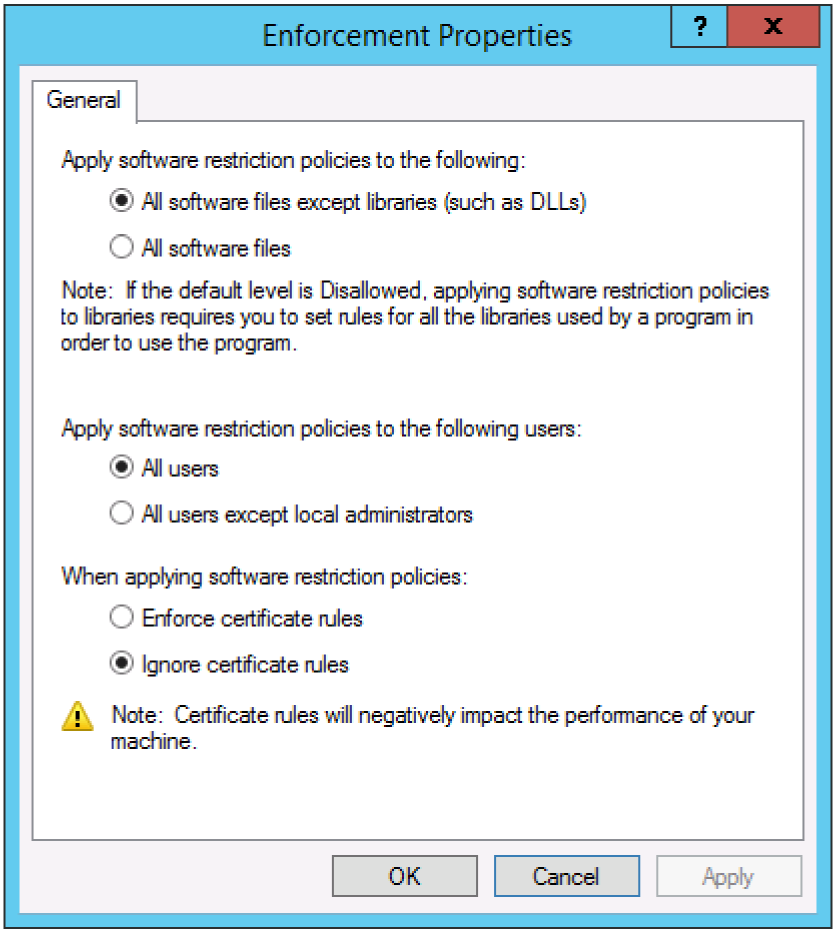

Apply software restriction policies to the following:

● All software files except libraries (such as DLLs)

○ All software files

Note: If the default level is Disallowed, applying software restriction policies to libraries requires you to set rules for all the libraries used by a program in order to use the program.

Apply software restriction policies to the following users:

● All users

○ All users except local administrators

When applying software restriction policies:

○ Enforce certificate rules

● Ignore certificate rules

⚠ Note: Certificate rules will negatively impact the performance of your machine.

Configure Windows Firewall

Note that this section on configuring Windows Firewall falls under the Create and manage Group Policy functional group. That means that most, if not all, or the Windows firewall related questions will be tied to Group Policy and not configuring or troubleshooting outside of Group Policy. Watch for questions concerning firewall settings in a GPO and familiarize yourself with the settings and implementation.

Configure rules for multiple profiles using Group Policy

There are 3 profiles that you can independently configure when configuring the Windows firewall:

- **Domain profile settings**. This profile is used for domain-joined computers when they are connected to a network and have connectivity to a domain controller. This is considered to be the least restrictive profile.
- **Private profile settings**. This profile is used for private networks such as home networks and business networks when computers are not joined to a domain. This is considered to be more restrictive than the domain profile but less restrictive than the public profile.
- **Public profile settings**. This profile is used in public places such as restaurants, airports, and libraries.

Starting with Windows 7 and Windows Server 2008 R2, a computer can operate under multiple profiles at the same time. Prior to Windows 7 and Windows Server 2008 R2, a computer could operate under a single profile. The screen capture below shows the configuration windows for configuring the different profiles in a GPO:

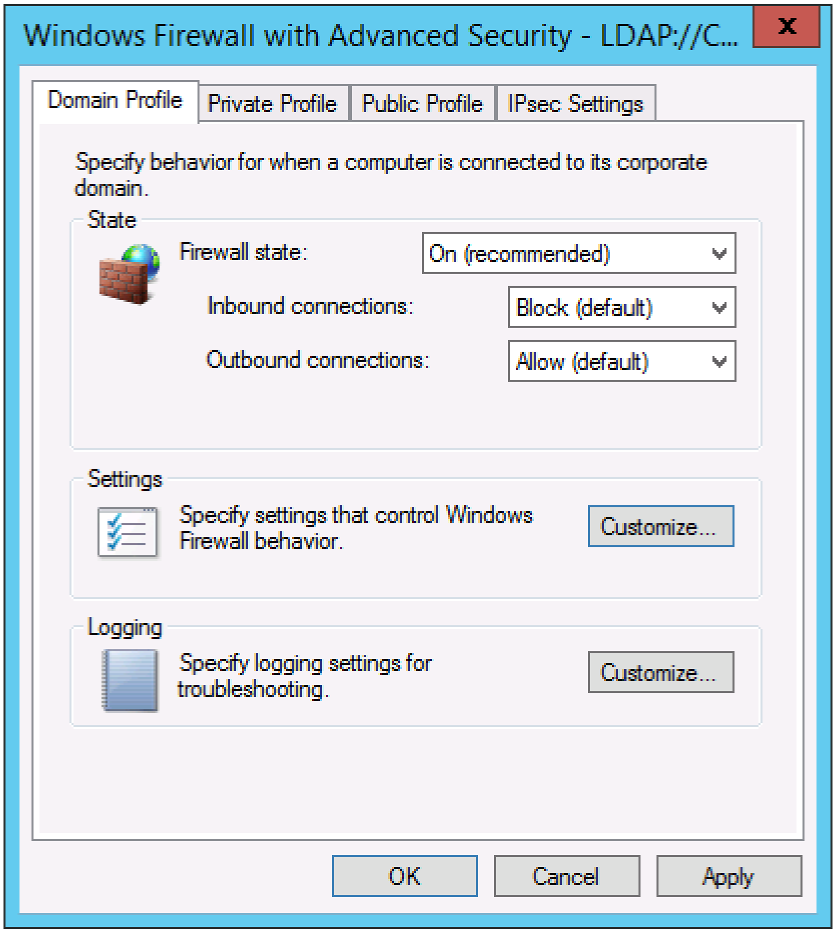

Windows Firewall with Advanced Security - LDAP://C...

Specify behavior for when a computer is connected to its corporate domain.

State

Firewall state: On (recommended)

Inbound connections: Block (default)

Outbound connections: Allow (default)

Settings

Specify settings that control Windows Firewall behavior.

Logging

Specify logging settings for troubleshooting.

In addition to turning the firewall state on or off and configuring connections per profile, you can also configure notifications and logging per profile. For example, if you wanted to configure a notification when inbound communication was blocked but you only wanted that for the public profile, you could configure that in the Public Profile tab. From an exam perspective, watch for scenarios that require minimizing the amount of logs captured or have unique requirements per profile. You can configure unique requirements per profile in a single GPO.

Configure connection security rules

Connection security rules are rules for how communication occurs over IPsec. You combine connection security rules with inbound and outbound rules to enable communication and dictate how the communication occurs. The screen capture below shows the options for creating new rules:

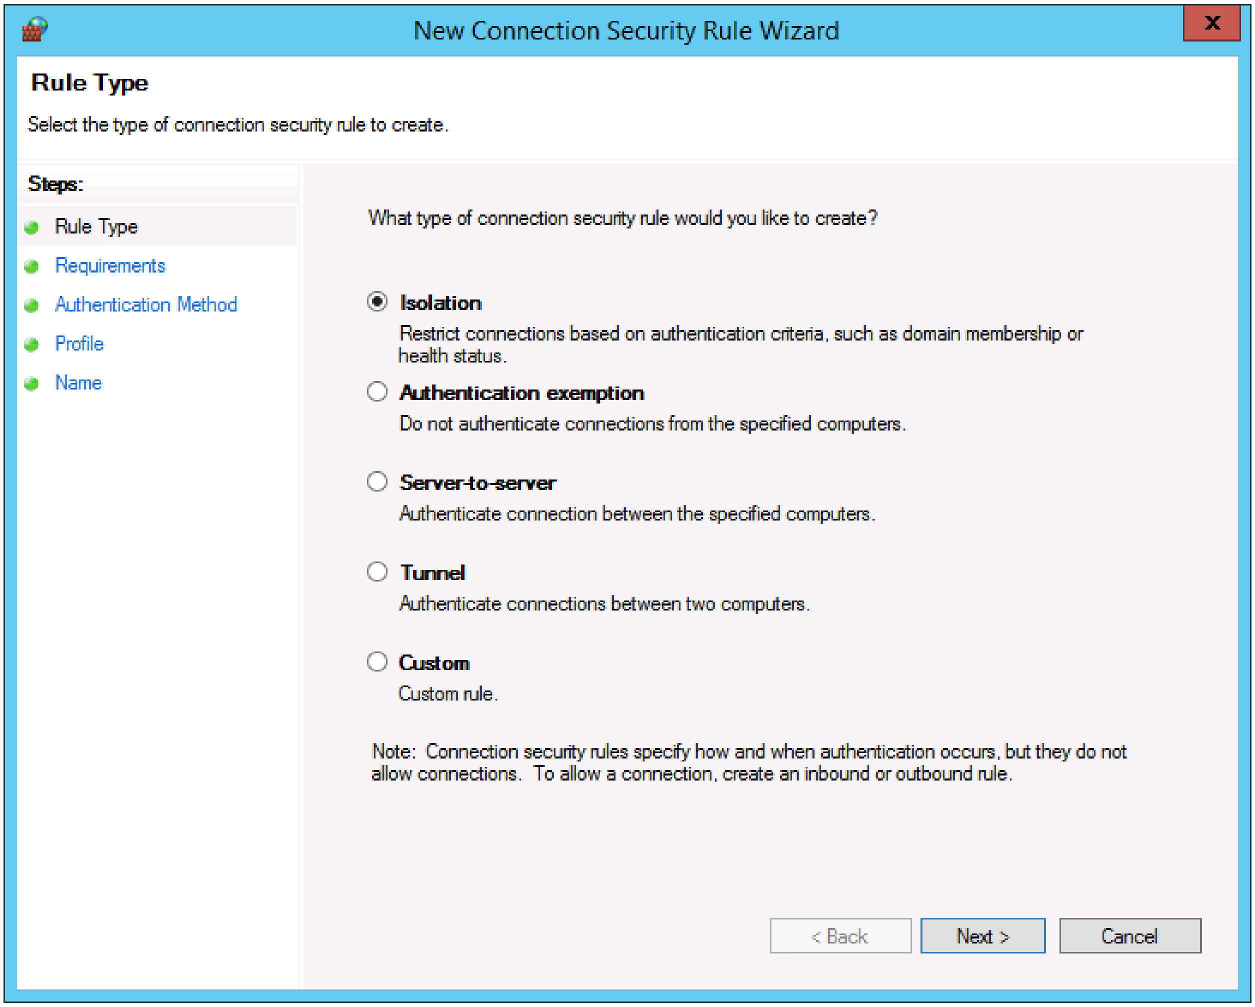

New Connection Security Rule Wizard

Rule Type

Select the type of connection security rule to create.

Steps:
- Rule Type
- Requirements
- Authentication Method
- Profile
- Name

What type of connection security rule would you like to create?

● **Isolation**
Restrict connections based on authentication criteria, such as domain membership or health status.

○ **Authentication exemption**
Do not authenticate connections from the specified computers.

○ **Server-to-server**
Authenticate connection between the specified computers.

○ **Tunnel**
Authenticate connections between two computers.

○ **Custom**
Custom rule.

Note: Connection security rules specify how and when authentication occurs, but they do not allow connections. To allow a connection, create an inbound or outbound rule.

Be aware of the following information for the exam:

- You can use connection security rules to ensure that communication between computers is encrypted.
- You can use connection security rules to ensure that computers authenticate each other before they begin communication.
- You can use IPsec between two computers without having the communication encrypted.
- You can use SSL certificates, NTLMv2, or Kerberos to authenticate computers. In a domain environment, Kerberos is the best choice because of the low administrative overhead. In an environment where there is a need to authenticate computers that run operating systems other than Windows, SSL certificates is the best choice because they are universally supported across most operating systems.

Configure Windows Firewall to allow or deny applications, scopes, ports, and users

You have a lot of flexibility when you configure the Windows firewall with a GPO. It is important to know the capabilities for the exam. The screen capture below shows the first window when you are adding a new inbound rule:

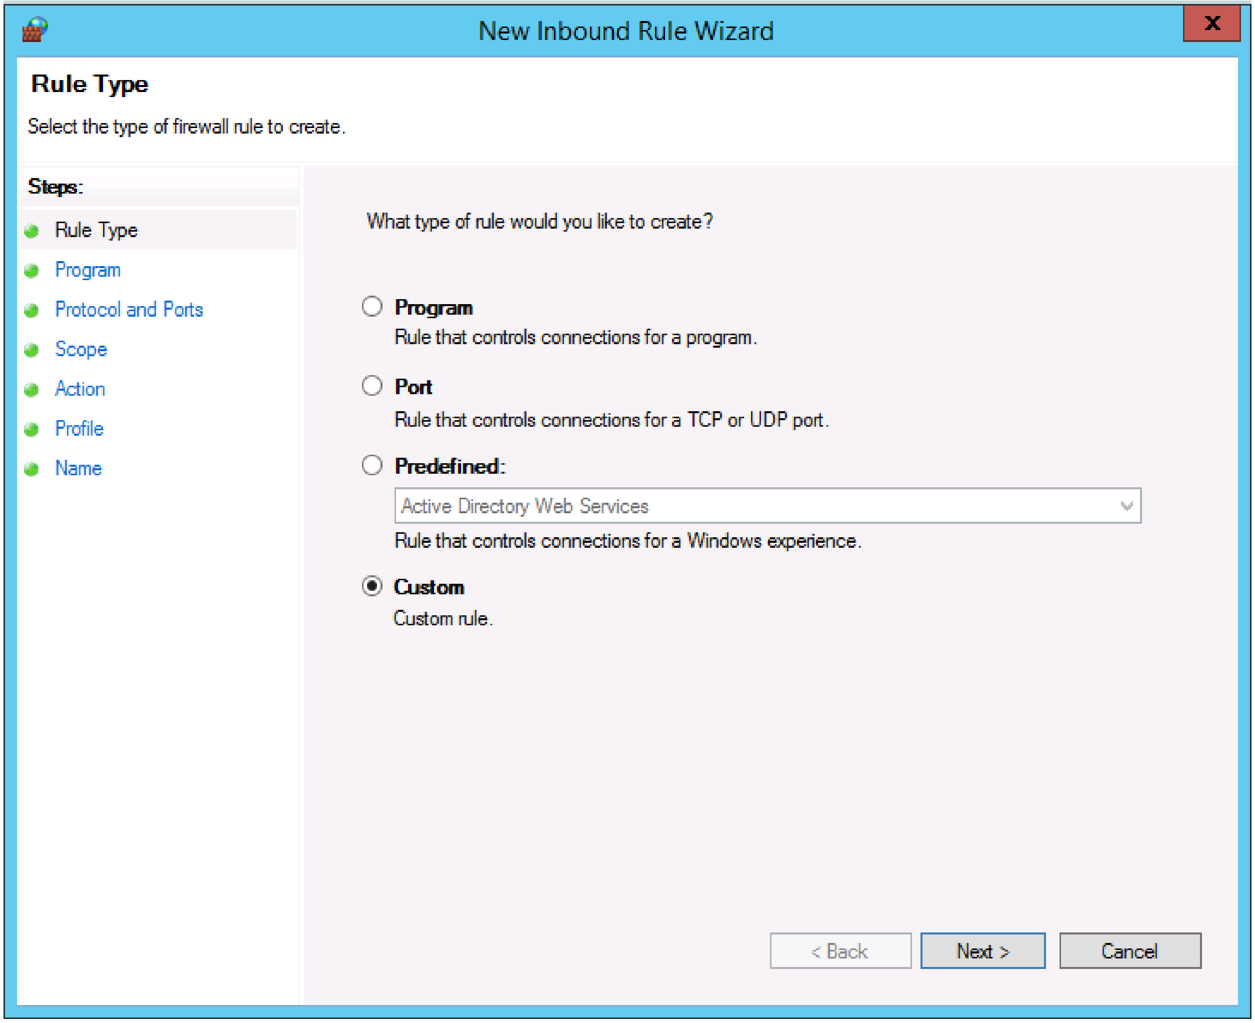

Rule Type

Select the type of firewall rule to create.

What type of rule would you like to create?

○ **Program**
Rule that controls connections for a program.

○ **Port**
Rule that controls connections for a TCP or UDP port.

○ **Predefined:**
Active Directory Web Services
Rule that controls connections for a Windows experience.

● **Custom**
Custom rule.

There are 4 rule types that you should know about:

- **Program**. This rule type points to an application. You point to the application's executable file, then specify whether you want to allow (or allow if a secured connection) or block the connection. You can also configure the rule for just one or more profiles.
- **Port**. This rule type specifies a port or group of ports. You can specify one port or a group of ports and then specify whether you want to allow (or allow if a secured connection) or block the connection.
- **Predefined**. This rule type specifies a known and existing service. You specify the service, such as SNMP traps, and then specify whether to allow (or allow if secured connection) or block the connection. You don't need to know the ports or programs which makes this rule the best for less experienced administrators. However, it may be compelling on the exam for scenarios that call for reducing administrative overhead.

- **Custom**. This rule type can be configured for programs, ports, or known services. With a custom rule, you can create very granular rules. For example, you can create a rule for iexplore.exe to allow only port 443 while using wireless and the public profile. On the exam, if there are very complex requirements for a firewall rule, the custom rule could be the route to go.

Rules can apply to users, to computers, or to both. But an important caveat exists for creating rules that only work for a specific user or group. You have to secure the connection to specify users. So, you could allow a connection if it is authenticated and integrity-protected with IPsec while specifying only a specific user is allowed to make the connection. For the exam, watch for scenarios that need a firewall rule to apply only to a user and then look for answer choices that have connection security.

Configure authenticated firewall exceptions

Firewall exceptions allow you to specify a user or group that won't be part of a firewall rule. For example, if you have a kiosk computer and you want to deny internet access to the browser, you could create such a rule and exclude administrators that might need internet access to service and maintain the kiosk computer. The screen capture below shows the configuration of specifying users and exceptions.

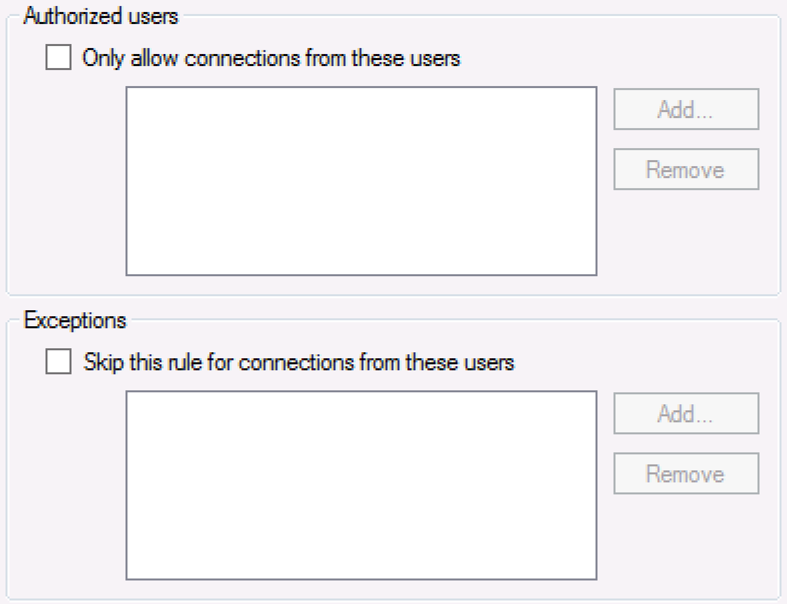

Authorized users

□ Only allow connections from these users

Exceptions

□ Skip this rule for connections from these users

Import and export settings

When you configure firewall settings in a GPO, you have the option to import settings to reduce the administrative overhead of manually creating settings. For the exam, know the following information:

- You can export firewall settings in an existing GPO to a .wfw file.
- When you import firewall settings from a .wfw file to a GPO, the settings overwrite the current settings.
- Windows firewall policy files (.wfw files) are not human-readable. They are binary files and cannot be modified.

Useful References

- Windows Server Auditing Free Quick Reference Guide
- SysAdmin Magazine May: Basic Rules of Windows Server Security
- How to Enable Video Recording of Changes in Your Windows Server
- How to Detect Who Created a Scheduled Task on Windows Server in Real Time
- Ten Simple Ways to Prevent Data Breaches in Windows Server 2012
- How to Detect Unauthorized Software Installation on Windows Server
- How to Detect Password Changes in Active Directory
- How to Disable Inactive User Accounts Using PowerShell
- Local Administrator Group Changes: Get Notified with PowerShell
- How to Detect Changes to Organizational Units and Groups in Active Directory
- Monitoring Event Logs with PowerShell
- Add Sensitive User Accounts to the Active Directory Protected Users Group\
- IT Audit & Compliance: Top Webinars to Attend

Complete Visibility of IT Infrastructure with Netwrix Auditor

Netwrix Auditor is an IT audit software that maximizes visibility of IT infrastructure changes and data access. The product provides actionable audit data about who changed what, when and where and who has access to what.

You can learn more about Netwrix Auditor and download a free 20-day trial.

“...best Active Directory/Group Policy product and Best Auditing/Compliance product 4 years in a row...”

“...auditing is generally a rather difficult task, especially if done manually. All of the many details you need to consider and remember are taken care of by Netwrix Auditor...”

About the Author

Brian Svidergol is focused on Microsoft infrastructure and cloud-based solutions around Windows, Active Directory, Exchange, System Center, virtualization, and MDOP. He holds numerous certifications including MCITP, MCSE, RHEL3, VCP, NCIE-SAN, MCT, MCSA, Microsoft Certified Solutions Expert: Server Infrastructure. Brian is an author of Microsoft Official Curriculum (MOC) course 6426C - Configuring and Troubleshooting Identity and Access Solutions with Windows Server 2008 Active Directory. He has worked on Microsoft certification exam development and related training content for several years. Also, he has co-authored the Active Directory Cookbook.

About Netwrix Corporation

Netwrix Corporation is the IT auditing company, providing software that maximizes visibility into who changed what, when, where and who has access to what. Over 6,000 customers worldwide rely on Netwrix to audit IT infrastructure changes and data access, prepare reports required for passing compliance audits and increase the efficiency of IT operations. Founded in 2006, Netwrix has more than 70 industry awards and named to the Inc. 5000 list and Deloitte Technology Fast 500.

Netwrix Corporation, 300
Spectrum Center Drive, Suite
820 Irvine, CA 92618

Regional offices:
New York, Atlanta, Columbus, London

netwrix.com/social

Toll-free: 888-638-9749

Int'l: +1 (949) 407-5125

EMEA: +44 (0) 203-318-0261