LabSim

TestOut: Windows Server Pro

Install and Configure

Video Transcripts

TestOut®

Module	Sections	Time	Total	HR:MM
1.0 Introduction				
	1.1 Windows as a Server	10		
	1.2 Windows Server 2008 R2 Interface Overview	25		
	1.3 Windows Server 2012 Interface Overview	25	60	1:00
2.0 Installation and Configuration				
	2.1 Installation	20		
	2.2 PowerShell	35		
	2.3 Server Roles	35		
	2.4 Server Core	35		
	2.5 Remote and Offline Servers	50		
	2.6 NIC Teaming	45		
	2.7 Traditional Storage	75		
	2.8 Storage Pools	50	345	5:45
3.0 Hyper-V				
	3.1 Virtual Machines	55		
	3.2 Virtual Machine Storage	55		
	3.3 Virtual Networks	60	170	2:50
4.0 Active Directory				
	4.1 Active Directory	20		
	4.2 Domain Controllers	60		
	4.3 Sites	10		
	4.4 Organizational Units	30		
	4.5 User Accounts	50		

8.4 Audit Policies	35		
8.5 User Rights Assignment	10		
8.6 Security Options	40		
8.7 Restricted Groups	15		
8.8 Windows Firewall Policies	30		
8.9 Application Restriction Policies	50		
8.10 Group Policy Preferences	10	315	5:15
9.0 Networking			
9.1 IPv4 Addressing	50		
9.2 IPv4 Custom Addressing	40		
9.3 IPv6 Addressing	40	130	2:10
10.0 DHCP			
10.1 DHCP Basics	40		
10.2 DHCP Exclusions and Reservations	25		
10.3 DHCP Centralization	40		
10.4 DHCP Troubleshooting	20	125	2:05
Windows Server Pro: Install and Configure Practice Exams			
Objective 1: Configure Windows Servers (6 simulation questions)	30		
Objective 2: Hyper-V (6 Simulation questions)	30		
Objective 3: Active Directory (10 simulation questions)	50		
Objective 4: DNS (7 simulation questions)	35		
Objective 5: File and Print Services (19 simulation questions)	95		
Objective 6: Group Policy (9 simulation questions)	45		

	Objective 7: Networking and DHCP (9 simulation questions)	45		
	Windows Server Pro: Install and Configure Certification Practice Exam (15 simulation questions)	75	405	6:45
Microsoft 70-410 Practice Exams				
	Objective 100: Install and Configure Servers (70 questions)	70		
	Objective 200: Configure Server Roles and Features (69 questions)	69		
	Objective 300: Configure Hyper-V (35 questions)	35		
	Objective 400: Deploy and Configure Core Network Services (83 questions)	83		
	Objective 500: Install and Administer Active Directory (58 questions)	58		
	Objective 600: Create and Manage Group Policy (61 questions)	61		
	Microsoft 70-410 Certification Practice Exam (60 questions)	60	436	7:16
		Total Time	**2806**	**46:46**

1.1 Windows as a Server

This course is designed to prepare you for the following certification exams:

- TestOut's Windows Server Pro: Install and Configure
- Microsoft's 70-410 exam, a requirement for:
  Microsoft Certified Solutions Associate (MCSA)
  Microsoft Solutions Expert (MCSE): Server Infrastructure
  Microsoft Solutions Master (MCSM): Directory Services

The Windows Server Pro: Install and Configure Certification is from our new line of TestOut Pro Certifications--certifications that measure not just what you know, but what you can do. The Windows Server Pro: Install and Configure Certification measures your ability to design, implement, configure, and manage a Windows network that incorporates Windows Server 2012.

The MCSA certification measures a primary set of Windows Server 2012 skills, relevant across multiple solution areas in a business environment.

The topics covered in this course are:

- Install and configure servers
- Configure Hyper-V
- Install and administer Active Directory
- Implement DNS
- Deploy and configure core network services
- Configure server roles and features
- Create and manage Group Policy
- Configure network settings
- Implement DHCP

1.1.1 Windows Server

Windows Server
0:00-0:05
Welcome to Windows Server 2012. I want to talk just a couple of minutes about what a server is.

Server
0:06-1:10
I think Client and Server are two of the most difficult terms to define in our industry, because they can mean so many different things. If you said to me, "Hey, Shad, I went out last night and I built a server", I don't necessarily know exactly what you mean. Maybe you built up some really awesome hardware that's intended never to be able to fail and put in redundant hard drives and multiple network cards. Maybe you just installed Windows Server 2012. Maybe you shared a file on your Windows 8 desktop. All of those things, technically, would be a server.
Really, the idea behind the server is, it's going to provide services to clients. The clients themselves can function as servers, but the idea behind the server is, it's a larger environment, and we're looking to centralize control over our resources. The idea behind networking is to share information and resources so we don't have to put a 3 TB drive in everybody's computer at their desktop.
The main themes that I'd like to have my students focus on really break down into three ideas. The first idea is, when you want to be a really efficient network administrator, you want to be as lazy as possible.

Efficiency
1:11-1:45
Most students laugh at that, but it's really true. In an ideal environment, I'm going to sit at my desk and probably not get up. I should join a gym because I would be sitting there most of the time. If I can't remote desktop into the server, if I can't control it from my desk, that's a bad day in my network. My boss is probably ringing my cell phone off the hook, the users are screaming, it's dangerous to walk in the halls--that type of thing.
As you go through your career, you always want to think of, "How can I work efficiently as opposed to harder?" Going along with that idea of efficiency, really, there's two main themes.

Centralized Security and Centralized Administration
1:46-2:44
They both revolve around centralization. With server, we're always looking for centralized security, which in our case is going to be Active Directory, a centralized database, and centralized administration. When I look at the computers in my network, I don't want to have to visit any individual computer. In fact, you might be surprised to realize, but the most difficult computers to support are computers of your people: in your friends, in your family, where they're very individualized.
In a network, I want everybody's computer to be the same. It's not their computer; it's the company's computer. We want centralized administration, that way, I can make changes on all of my computers; that they remain identically configured and configured according to company standard. When we get in later on in the course, the centralized administration piece is really going to be Group Policy. We're focusing on a server, we're focusing on centralization, working efficiently, and having everything be standardized across the environment.

1.1.2 Windows Server Facts

A server is designed to manage access to centralized resources or services in a network. Servers are often identified by the services they provide, such as:

- File servers
- Print servers
- Email servers
- Web servers
- Proxy servers
- FTP servers

Microsoft Server operating systems are designed to facilitate:

- Centralized network administration
- Centralized network security
- Standardized network deployment
- Standardized security implementation

Windows Server 2012 introduces new and improved features to enhance traditional services as well as to provide for a modern IT organization's evolving needs to support virtualization and cloud-based applications and services. Enhanced functionality includes:

- Server Manager facilitates managing multiple local and remote servers from one management console.
- Hyper-V Server provides hypervisor-based virtualization.
- Hyper-V Manager centralizes administration of virtual machines and virtual networks.
- Windows PowerShell 3.0 provides comprehensive management capabilities from the command line and a Robust Session Connectivity feature to protect against damage caused by disconnection.
- Storage spaces provide virtual storage that can be dynamically managed and eliminates the need for such tasks as repartitioning drives, resizing volumes, and backing up data into order to repartition.
- Active Directory enhancements include domain controller cloning, Dynamic Access Control for easier access authorization, and automatic generation of PowerShell commands by the graphical user interface.
- IP Address Management (IPAM) provides network administrators a single console from which they can view and manage the IP addresses of an entire enterprise.

1.2 Windows Server 2008 R2 Interface Overview

As you study this section, answer the following questions:

- Which command can you use to display the ICT?
- How should you access Windows PowerShell to ensure that you have the PowerShell commands for the roles and services you have installed?
- What tasks do you perform to customize a server?
- What are key differences between Windows Server 2008 and Windows Server 2008 R2?
- What is the difference between roles, role services, and features?

After finishing this section, you should be able to complete the following tasks:

- View available Administrative Tools from the Start menu.
- Open Administrative Tools within Server Manager.
- Execute a command from a command prompt.
- Access the Control Panel.
- Open the Network and Sharing Center.
- Open Computer Management within Administrative Tools.
- Open Windows Explorer and browse the Windows Server 2008 R2 file structure.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  Navigate Server Interfaces
  - Navigate the Windows Server 2008 R2 Interface

This section covers the following 70-410 exam objective:

- 100 Install and Configure Servers.

1.2.1 Windows Server 2008 R2 User Interface

Windows Server 2008 R2 User Interface

0:00-0:51

Before we talk about Windows Server 2012, let's talk a little bit about Windows Server 2008 R2, which is the operating system before Windows Server 2012. The interface for Windows Server 2008 R2 is very similar to Windows 7. And one thing you might want to be aware of is that the server operating systems always come out with a generation of client operating systems that have similar interfaces and what's under the hood, so to speak, is pretty much the same. In fact, Microsoft builds in limitations into the clients so that you can't just use it as a server and get out of paying for an actual version of the server.

The 2008 R2 interface has a streamline installation. So when you are looking at that interface, there is going to be a few things that you will be aware of. When you first initially install that server, the first thing you are going to see is called the Initial Configuration Task or ICT Window.

Initial Configuration Task or ICT Window

0:52-1:55

The ICT is to initially configure the server. Way back in Windows Server 2003, when you installed it, it would stop periodically and ask you questions. What do you want to do for networking? What do you want to name this server? What do you want the administrator's password to be?

Starting with Windows Server 2008, which went with the Vista Client operating system, they streamlined the installation so that once you choose the hard drive where you are going to install it, you're not asked any other questions until the installation is completely finished. You give the administrator password, and then that's about it. Well, all those questions they used to ask you during installation still need to be set up. Most important is the name of the server and the time zone. Time is particularly important in a network.

So, the ICT window is going to pop up to make sure you make these changes. You set the correct time zone, you set the name of the server. You set up whether or not the firewall is going to be on, whether or not you are going to use Windows Update.

Close the ICT Window

1:56-2:22

If at any time you close the ICT, and it'll keep popping up until you click the checkbox in the lower right-hand corner that says, never show me this ever again. If for some reason you check that and you need to get it back. you can go into the Start button. Click on Start, and in the search box, you can type OOBE.

Now, that having been said, there is nothing that can be done in the ICT that can't be done in the Windows Server 2008 R2 interface.

Windows Server 2008 R2 Interface

2:23-3:16

Mainly for working with 2008 R2, you either work in Server Manager, which was intended to be a centralized administration program, or, I'm kind of old school, I like to use the tools that are in the Administrative Tools Menu, Start, Administrative Tools, whichever way you prefer to work with server--those are generally what you are going to use to configure the server itself.

It looks a lot like Vista and Windows 7. It's got a Start button but it doesn't have what they call the Desktop Experience. So it's not going to be the round Start button that just has the Windows logo on it. It's going to be, kind of, the old school button with the word Start on it. The graphics are mostly grey tone, because that doesn't put a lot of burden on the video processor of the server. We don't need those translucent windows or Aero experience on a server. In an ideal world, I won't interact with the server very much.

Working with the Server

3:17-3:55

A lot of students say to me, "Shad, I can't wait to work with the server," and I say," I hope I don't have to work with the servers too much". I'd like to get them configured, up and running, put them in the server room, should be a locked spot, and then they'll stay locked in there, and the only time I have to visit them is if I have to do something that has to be done physically with the hardware. If you are still using backup tapes, maybe you have to switch backup tapes. Generally speaking, I don't want to visit the server very often, at all.

So, there is not much there by way of graphics, but a very similar interface to all of the generations that we've seen with the Windows products, and we will be sticking in Server Manager and in the Administrative Tools menu.

1.2.2 Using the Windows Server 2008 R2 User Interface

Using the Windows 2008 R2 User Interface
0:00-0:09
In this demo, we're going to take a look at the Windows Server 2008 R2 Interface.
You can see here, the Initial Configuration Tasks window is open.

Basics for Every Server
0:10-0:45
Once you install 2008 R2, it will prompt you for the administrator password, and then when you first log in, it's going to open up this ICT window that prompts you to do the basic things that should be done for every server.
We have Activate Windows, Set the time zone, Configure networking, Provide the computer name and domain, set up automatic updates, so on and so forth. Anything that you can do in this window can be done elsewhere in the operating system. It's really just intended as a check list for what you should do with a new server.

Removing and Adding the ICT Window
0:46-1:10
If you don't want to see this window anymore, you can come down in the bottom left hand corner and check "Do not show this window at logon". When I hit Close, next time I log in, I'm not going to see it. If you ever want it back, all you have to do is go into the Start button. Right in the search box, you type oobe, the window comes back up, and you can uncheck it.

Opening Server Manager
1:11-1:26
Once we close the ICT window, normally what it will do is send this into Server Manager. If it doesn't, I just like to come right down here to the right of the Start button, click the Server Manager button, and that will open up Server Manager.

Server Manager
1:27-1:48
Server Manager is intended to be kind of a main application. I can manage the Roles that are installed on the computer. We can see that this particular server has quite a number of roles installed; Certificate Services, Domain Services, Application Server. Then, I can click on any one of these individual nodes, and I should get some options for managing that role.

Roles
1:49-1:55
Roles are main functions of the server and you can see right here on roles, I can Add Roles. I can Remove Roles.

Features
1:56-2:16
Features are more supporting functions, so not anything that the server is mainly doing. Here, it's got the Desktop Experience, so that it could look like a Windows 7 machine. Or it's got some Group Policy Management or Network Load Balancing. These aren't the primary functions of the server, but they're features that are supporting it.

Subset
2:17-3:16
The only thing you really have to know for Server 2008 R2 is, if it's not a Role, it's a Feature, and if it's a subset of a role, once the role is installed, you have to modify the role services. For example, if I scroll down in Roles here, you can see that in my Certificate Services, it only installed the Certification Authority. These are the other Role Services that are available with that role, and I would have to come in here and say Add Role Services in order to have these things be installed.
When you first install the role, it prompts you for which role services you want. Once it's already installed, you have to find that section in roles and do the Add Role Services. That's pretty much it for the 2008 R2. If you don't like Server Manager, you can get to individual snap-ins by going through the Start button, and all of the snap-ins will be added to Administrative Tools.

PowerShell
3:17-3:46
With PowerShell, the other thing I'll point out to you while we're here, you can see that up here, we have an Active Directory Module for Windows PowerShell. That would open up PowerShell with all the Active Directory commands already imported. The PowerShell that's on the task bar, down here to the right of the Start button, is just a generic PowerShell. If I open it that way, I wouldn't have any of the Active Directory commands. I'd have to do an import module to make those available.
That's just a brief demo of the Windows Server 2008 R2 Interface.

1.2.3 Windows 2008 R2 User Interface Facts

The main differences between Windows Server 2008 and Windows Server 2008 R2 are:

- Windows Server 2008 is the same codebase bits as Vista. It is available in 32-bit and 64-bit versions.
- Windows Server 2008 R2 is the same codebase bits as the Windows 7 64-bit version. It is available only in the 64-bit version.

The tools used to administer Windows Server 2008 R2 are described in the following table:

Tool	Description

Initial Configuration Tasks (ICT)	After you install Windows Server 2008 R2 and enter your password, the Initial Configuration Tasks (ICT) screen displays. The ICT functions as a checklist for setting up a server. The information you enter to set up the server is categorized into the following areas:
	• Provide Computer Information
	  • Activate Windows
	  Set the time zone
	  Configure networking
	  Provide computer name and domain
	• Update the Server
	  • Enable automatic updating and feedback
	  Download and install updates
	• Customize This Server
	  • Add roles
	  Add features
	  Enable Remote Desktop
	  Configure Windows Firewall
	All of the tasks available in the ICT can be completed in other areas of the Windows Server 2008 R2 interface. To display the ICT, enter **oobe** in the Search Box.

Server Manager	Server Manager provides a single source for managing a server's system information. Server Manager eliminates the requirement that administrators run the Security Configuration Wizard before deploying servers. Server Manager: • Manages roles and features installed on the server. • Displays server status • Identifies problems with server role configuration

	Server Manager replaces several features included with Windows Server 2003, including: • Manage Your Server • Configure Your Server • Add or Remove Windows Components By default, server roles are configured with recommended security settings and are ready to deploy as soon as they are installed and properly configured.
Administrative Tools	You can access Administrative Tools from the Start menu. All of the snap-ins available in Server Manager are available in Administrative Tools. The Active Directory Module for Windows PowerShell available through Administrative Tools has all of the PowerShell Active Directory commands imported. PowerShell available on the Taskbar is a generic version of PowerShell and will not have the Active Directory commands available.

The following are editions of the Windows Server 2008 operating system:

- Windows Server 2008 Standard (32-bit and x64 versions)
- Windows Server 2008 Enterprise (32-bit and x64 versions)
- Windows Server 2008 Datacenter (32-bit and x64 versions)
- Windows Web Server 2008 (32-bit and x64 versions)
- Windows Small Business Server 2008 Standard edition
- Windows Small Business Server 2008 Premium edition (32-bit and x64 versions)
- Windows Essential Business Server 2008 Standard and Premium editions

The following are editions of the Windows Server 2008 R2 operating system:

- Windows Server 2008 R2 Standard
- Windows Server 2008 R2 Foundation
- Windows Server 2008 R2 Enterprise
- Windows Server 2008 Datacenter R2
- Windows Web Server 2008 R2

1.3 Windows Server 2012 Interface Overview

As you study this section, answer the following questions:

- Which Windows Server 2012 editions provide all features of Windows Server 2012?
- How do you access the Charms menu?
- What are two ways can you access the Start screen?
- How do you add tools to the Tools menu?
- What happens when two commands are piped?

After finishing this section, you should be able to complete the following tasks:

- Open Administrative Tools within Server Manager.
- Access the Start screen.
- Open an application from the Start screen.
- Run a command from a command prompt.
- Access the Control Panel.
- Open the Network and Sharing Center.
- Open Computer Management within Administrative Tools.
- Open File Explorer and browse the Windows Server 2012 file structure.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  - Navigate Server Interfaces
    - Navigate the Windows Server 2012 Interface

This section covers the following 70-410 exam objective:

- 100 Install and Configure Servers.

1.3.1 Windows Server 2012 User Interface

Windows Server 2012 User Interface
0:00-0:06
One of the most striking things about Windows Server 2012 is the way Microsoft has redesigned the interface.

No Start Button
0:07-0:26
The first obvious thing that you're going to notice is, there won't be a Start button anymore, which was kind of a shock to me. I've been working with the Start button since Windows 95. Suddenly, it's gone, and for a while, I was a little bit lost.
There is no Start button. As soon you install server, Server Manager will open.

Server Manager
0:27-1:01
In 2008 R2, Server Manager was supposed to be the central command spot for Windows. It's much more so in Server 2012. Not only does it allow you to manage the local server to do all of the functions that you need to do within that server, it also allows you to remotely manage most of the other servers in your environment. It can handle about up to 100 servers comfortably. It's really is a one-stop shop for managing your network, unless you have a really Enterprise network with more than 100 servers.

Tools Menu
1:02-1:09
It also has a Tools menu that allows you to get access to all of the administrative snap-ins.

Start Menu
1:10-2:05
You can still go through the Start menu. The Start menu exists; we just don't have a Start button. You can go through the Start menu and go into Administrative Tools, if you want to, but it actually might end up being easier just to go to Tools inside of Server Manager and pull up the tool that you want.
You are going to see that this version of the operating system is more keyboard-- intensive than I've seen in a long time. But if you want to launch the Start menu you will hit the Windows key on the keyboard and it's going to open up with what they are calling the sort of the metro interface. You've got tiles there. It's not as interactive as Windows 8 because you won't have live tiles, but as we install roles into the server the different tools will be added to the Start menu that if you hit that Windows key most of your main tools will be right there.
Again, it's probably going to end up being more efficient to stay in Server Manager but certainly the Start menu is there.

Search
2:06-3:38
One of the big differences between having a Start menu, but not a Start button is not having the search box. If you are used to Windows Server 2008 R2 or Vista or Windows 7, any of those, you are used to clicking on the Start button and you see the search box right there, you start typing, whatever you need comes up in the Menu.
With Server 2012, there is a search function, but it's not so easy to get to. What you are going to have to do is hover your mouse either in the far upper right-hand corner of the screen or the far lower right-hand corner of the screen. That's going to pop out a bar with a search button on it. It's also where you get to your power settings so that you can shut down or reboot the server, but it takes a little bit of getting used to. You have to get into the framework of, okay, if what I see isn't exactly in front of me, I'm going to the upper right-hand corner, or I'm going to lower right-hand corner, then I'm going to be able to go into search and pull up what I need. But, the idea was that everything would be in Server Manager, and you wouldn't have to go into search very often at all. And hopefully it really does work out like that for you, but the interface will take some getting used to.
Initially, I will admit, I was very reluctant to move into a situation with no Start button, but I found that after a couple of weeks of working with it, you do get used to it, and it's not as bad as it was. There is a shift in Microsoft's thinking much more towards the keyboard.

Command Line vs. GUI
3:39-4:41
As we get into the course later on, we are going to see things like PowerShell; we are going to talk about command-line scripting. The philosophy being, that the command-line has always been probably more powerful than the GUI. The criticism with the Graphical User Interface is that it's bulky. It adds a video burden to the computer, so the computer has to display the graphics of the GUI. And because you have more code going on there, you have more potential for needing patches or having security holes. Later on in the course, we'll talk about Server Core, which doesn't have a GUI at all.

There is a trend towards working in the command-line, working with the keyboard, and it starts right out in the immediate interface, where you have to hit the Windows key on the keyboard in order to open the Start menu, or you've got to hover down there in order to get to the Search Menu.
I think once you experiment with it a little bit and get used to working with Server 2012, it won't be quite as bad as it might initially feel, and you may even come to like it.

1.3.2 Using the Windows Server 2012 User Interface

Using the Windows Server 2012 User Interface
0:00-0:12
This demo is just going to be a little tour of the Windows Server 2012 Interface. Once you install Server 2012 and you set your password, it will take you directly into Server Manager.

Server Manager
0:13-0:23
Server Manager is intended to be the main management utility. I think you could probably do almost everything you need to do and never leave Server Manager.

Dashboard
0:24-0:54
It takes us in on the dashboard, and you can see right here, we've got links for adding roles and features. Roles are main functions of the server, features are supporting functions of the server. I can also "Add other servers to manage" or "Create a server group", which we'll cover in a different demo. It also gives me a little snapshot down here of File and Storage Services, the Local Server, All Servers. If something was wrong, these would be red. Sometimes when it's booting, it will go red, and then it will turn to green afterwards.

Local Server
0:55-1:23
In Server 2008 R2, we had the Initial Configuration Task window. If I click on Local Server, you're going to see something very similar to that. After you first install the server, you want to come into Local Server, set up your Computer name, Windows Firewall, Remote Desktop. Most important is Computer name, Time zone, and Windows Updates, definitely at least set up all of those. Then you can go in and manage it.

Start Menu
1:24-2:18
Usually the most obvious difference with Server 2012 is the fact that we don't have a Start button. I think that's something that makes a lot of people feel uncomfortable. I know it made me feel uncomfortable. My first icon here is the icon on the task bar for Server Manager.
A couple of ways to get into the Start menu, because the Start menu is still there: One way is to hover the mouse in the way bottom left-hand corner. It pops up this little blue square, and then if we click on it, I'll see the Start menu. As I add roles to the server, it will add the snap-ins to the Start menu so that they would be right here.
The other way to get to the Start menu is just to hit the Windows key on the keyboard. That would take me right in as well. To get out of it, I can either click on Desktop, or I can hit Escape, and that will take me right back to my desktop.

Search Dialogue Box
2:19-3:05
The other thing I miss without having my Start button is the Search dialog box, which I did love. We still have it, but you've got to get used to hovering, so two ways to get it. You can hover your mouse in the upper right-hand corner or the lower right-hand corner. That pulls up this little bar here and then I can come down, I can get to my Start menu from here, I also can get into Search, where I can search for whatever I want.
As I put in text into the search box, it will tell me if there are Applications that match, Settings that match, or Files that match, and then I can click on whatever I need.
The other thing that's on this bar here is Settings.

Settings
3:06-3:42
We're used to going to the Start menu to shut down the server. Here in Settings, this is where I get my Power. I can go in and I can shut it down, or I can Restart it. I also have an onscreen Keyboard. Any Notifications will show up here. Then I can click on my network connections if I need to, as well. I've also got sound, and this would be the brightness of the display. There are some settings in there that you can configure as well.
That's really just a brief tour of the 2012 Interface.

Summary
3:43-4:41
You also still have the notification area down here, the bottom right-hand corner, where the clock and the date is, and some of the icons that you can use to get in there. It's a little bit different than Windows Server ... anything before 2012. Once you get into the habit of hovering in the corners and using the keyboard, it's really not a big deal to make an adjustment in how you do your work and get into the habit of launching the search or launching the Start menu.

Again, Server Manager is intended to be the main area. The beauty of Server Manager is this Tools menu. If I add in any roles, not only is the snap-in added to the Start menu, it's also added to this Tools menu. Conceivably, I could run the entire server. I might never need to to into search or into the Start menu, and be able to do everything I need to do just from within Server Manager.

1.3.3 Using the Windows Server 2012 R2 User Interface

Using the Windows Server 2012 R2 User Interface

0:00-0:19

You can see that Windows Server 2012 R2 is slightly different but not dramatically different than Windows Server 2012. We still have the Dashboard, the Local Server, and the All Servers, as well as the links for any rules that I've installed on my server. I still have the Tools menu and the Manage menu and my notifications.

Using the Start Button

0:20-0:47

The biggest difference is the addition of the Start button in the bottom left hand corner. I can click on the Start to take me to the Start menu or I can go ahead and right click Start and that will take me into Task Manager control panel or a bunch of cool utilities. It's a nice quick way to get around. I can even click shut down or sign out from here and not have to go into the Start menu and click my user name in order to sign out. Other than that, it's pretty much the same.

Using Search

0:48-1:16

The other thing that you might notice is when you go into search--if I hover down in the bottom right hand corner and I go into the Charm bar and click Search, it searches everywhere. In 2012, it used to categorize things into files, applications and you had to click in each of the categories to find what you were looking for. Here, it's going to show you everything. If I search for temp, it will search with anything with temp in it anywhere in the computer, and even sometimes things on the web.

Summary

1:17-1:23

A nice little bit of addition to the look and feel of the Windows Server 2012 R2, but again, not dramatically different from 2012.

1.3.4 Windows 2012 User Interface Facts

Windows Server 2012 uses Server Manager to manage networks and multiple remote servers from a single administration console. You can easily configure and manage remote servers using Server Manager or Windows PowerShell.

The following table describes the tools used to administer Windows Server 2012:

Tool	Description
Server Manager	After installing Windows Server 2012, Server Manager set up the server using the following steps: 1. Configure this local server. 2. Add roles and features. 3. Add other servers to manage. 4. Create a server group. Options available from the Manage menu are: • Add Roles and Features • Remove Roles and Features • Add Servers • Create Server Group • Server Manager Properties Tools for server management can be accessed through the Tools menu. Tools are automatically added to the Tools menu when additional server roles and services are installed. Roles and Server Groups display the status of each server and role. The status refresh time can be configured in Manage > Server Manager Properties. A server, server group, or role displayed in red indicates a problem. The following information displays for the local server, all servers, and installed roles: • Manageability • Events • Services (if applicable) • Performance • BPA results
PowerShell	Windows PowerShell is a command-line shell scripting language that allows you to administer, maintain, configure, and develop new features for Windows Server 2012. Designed especially for system administration, Windows PowerShell uses *cmdlets* to control and automate the administration of the Windows operating system and applications that run on Windows. PowerShell:

	• Is built on the .NET Framework. • Automates administrative tasks. • Provides access to data stores, such as the registry and certificate store, in the same way the file system is accessed. • Uses specialized, built-in PowerShell commands known as cmdlets that: Allow you to manage a computer from a command line. Use a verb and a noun separated by a hyphen. For example, **Get-Help**, **Get-Process**, and **Start-Service**. Can execute single commands or large scripts. Allow stringing together the actions of two or more cmdlets, known as *pipelining* (also called *piping*). In pipelining, output from the first cmdlet is fed into the second cmdlet and so on. PowerShell provides help for each cmdlet through the **Get-Help [cmdletname]** cmdlet. You can use the **Update-Help** cmdlet to download and install the newest versions of help topics for modules installed on your computer.

Windows Server 2012 uses the Metro interface introduced in Windows 8. The following table identifies interface elements used to navigate Windows Server 2012:

Interface Element	Description
Start screen	Tiles on the Start screen allow you to access Computer, Control Panel, Server Manager, and the desktop. You can pin tiles on the Start screen for the following items installed on your computer: • Desktop applications • Apps • Snap-ins You can access the Start screen by moving the mouse pointer to the lower-left corner or by pressing the Windows logo key. You can return to the desktop by pressing the ESC key or clicking the desktop tile.
Charms menu	Like Windows 8, Windows Server 2012 has a Charms menu. When you move the mouse pointer to the upper-right or lower-right corner of the screen, the Charms menu displays. You have the following options: • **Search** locates items matching searched keywords and displays the results below the Search box. • **Start** returns to the start screen. • **Settings** is divided into two parts: On the top of the panel you can access: ▪ Desktop

	▪ Control Panel ▪ Personalization ▪ Server info ▪ Help At the bottom of the panel you can access: ▪ Network ▪ Sound ▪ Screen brightness ▪ Notifications ▪ Power options ▪ Keyboard options

Windows Server 2012 has four editions:

- Windows Server 2012 Foundation
- Windows Server 2012 Essentials
- Windows Server 2012 Standard
- Windows Server 2012 Datacenter

Both Windows Server 2012 Standard and Datacenter editions allow an organization to use all Windows Server features. The main difference in the editions is the use rights for virtualization.

Windows Server 2012 R2

Windows Server 2012 R2 introduces several key changes to the graphical user interface, which are identified in the following table:

Interface Element	Description
Start button	The Start button has been added to the Taskbar in the same location where it was found in earlier versions of Windows Server. However, the functionality of the Start button in Windows Server 2012 R2 has changed in the following ways: • Clicking the Start button switches the system from the desktop environment to the Metro environment and displays the Start screen. • Right-clicking the Start button displays a pop-up menu with links to the following: Programs and Features Power Options Event Viewer System Device Manager Network Connections Disk Management Computer Management Command Prompt (standard or elevated permissions)

	Task Manager Control Panel File Explorer Search Run Shut down, sign out, or restart Desktop
Search charm	In Windows Server 2012, the Search charm presented search results sorted into categories, such as files or applications. Each category had to be manually expanded to view the search results. In Windows Server 2012 R2, the Search charm searches everywhere by default, including the Internet.
Boot to screen	By default, Windows Server 2012 booted to the Start screen. Windows Server 2012 R2 boots to the desktop by default. This behavior can be customized by completing the following steps: 1. In the desktop environment, right-click the Taskbar and select **Properties**. 2. Click the Navigation tab. 3. Under Start screen, either select or clear the **When I sign in or close all apps on a screen, go to the desktop instead of start** checkbox. 4. Click **OK**.

Windows Server 2012 R2 is available in four editions:

- Windows Server 2012 Foundation
- Windows Server 2012 Essentials
- Windows Server 2012 Standard
- Windows Server 2012 Datacenter

2.1 Installation

As you study this section, answer the following questions:

- Which editions of Windows Server 2012 support a Server Core installation?
- What is the system volume free space requirement for Windows Server 2012?
- If you are currently running Windows Web Server 2008 R2, what is your upgrade path?
- How many virtual instances are allowed on each Windows Server 2012 edition?
- What is the difference between a full installation of Windows Server 2012 and a Server Core installation?
- How does a Server Core installation in Windows Server 2012 differ from a Server Core installation in previous versions of Windows Server?
- Can a full installation of Windows Server 2012 be converted to a Server Core installation?

After finishing this section, you should be able to complete the following tasks:

- Plan a Windows Server 2012 installation.
- Install Windows Server 2012.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.

This section covers the following 70-410 exam objective:

- 101 Install servers.
  This objective may include but is not limited to:
  Plan for a server installation
  Plan for a server upgrade

2.1.1 Server Installation and Upgrade

Server Installation and Upgrade
0:04-0:14
Let's talk about Windows Server 2012 installation and upgrades. Back in Windows Server 2008 and 2008 R2, they really streamlined the installation process, so that you didn't have to answer very many questions before you actually got into the operating system.

Use of Product Key
0:15-1:10
One of the first differences we're going to see on the installation here is that you will have to put in a product key before the server will install. It used to be, you could install without a product key, and add the product key later. Here, we really can't do that. There are eval versions that you can get, and those can be converted into retail versions after the fact, but if you're installing a production server, the best thing is to have the serial code--the correct serial code--and go through the installation providing that code. That will install the appropriate version. Beginning in Windows Server 2008, all the versions of Windows Server should be on the same DVD. Windows knows which version to install based on the product key that you use. Once the installation is complete, then you'll go into the operating system and configure it the way you need it to be configured. First, we're going to go over the editions of Windows Server 2012. Then, we'll talk about the minimum requirements and some of the upgrade paths if you are working with servers that need to be upgraded.

Versions of Windows Server 2012.
1:11-1:12
Let's look at our different editions.

Foundation
1:13-2:02
Let's start by talking about foundation. This is a scaled down version of the operating system. It supports just one CPU socket and 32 GB of RAM. It might seem like a lot to you if you're used to dealing with client workstations, but for a server that might not actually be that much RAM. Certainly you can put some RAM in there. Here's one of the biggest limitations: only up to 15 simultaneous LAN connections. So this would be a very small environment, indeed. It supports up to 50 RRAS connections, up to 20 simultaneous remote desktop connections, if you're using remote desktop. But it doesn't give us our Server Core installation. It does not support Hyper-V and does not support server clustering. This is a very limited version of Windows Server 2012. It would be for small companies that don't have a lot of budget. Adequate for a small server that's going to be a file server or something like that, but again, a very small company.

Essentials
2:03-2:36
You could also get Essentials. This is a little bit better, you've got two CPU sockets, 64 GB of RAM, 25 simultaneous LAN connections. It is still a very small network. They have much more people coming in over RRAS and Remote Desktop, but again, no Server Core, no Server Clustering. This one's a little bit different. It can be installed into a Hyper-V guest operating system. You can run it as a virtual machine, but it cannot be hosting Hyper-V virtual machines. Again, this is a limited version of the operating system. These would be for small companies.

Standard
2:37-3:48
Most companies are going to have either Standard or Datacenter. Here, you can kind of see what I was talking about-- 1 CPU socket and 32 GB of RAM being kind of small.
Even with Standard, we get a potentially 64 CPU sockets, 4 TB of RAM. I would love to have a computer with 4 TB of RAM, unlimited connections, Server Core. Now, here's the cool thing, you can have up to 2 Hyper-V guest operating systems using the same license key as the host, assuming the host is only running Hyper-V. If you purchased 2012 Standard, you're really buying two Windows Server operating systems. Assuming that you install this on the host as Hyper-V, it only runs Hyper-V. Then, you create your server as guest operating system. Of course, it certainly can be installed in the Hyper-V as a guest, and it supports server clustering with up to 64 nodes per cluster. I don't want to do a whole section on clustering, but essentially clustering in a nutshell is allowing multiple servers to act as if they were one server. There's probably not one server that does Yahoo.com; it might be a cluster.
Here we can up to 64 nodes, which is a big improvement over previous versions of the operating system that limited it to more like 16. This is a huge difference coming into 2012.

Datacenter
3:49-4:24
Datacenter would be much more expensive and has much more capability. We still have our 64 CPU sockets and 4 TB, unlimited connections, supporting Server Core. Here we get unlimited Hyper-V guest operating systems using the

same license key as the host, the same 64 clustering. If you purchase Datacenter, conceivably you can install an unlimited number of Hyper-V guest operating systems. The key is that they'd have to be all on one box. You're going to have to buy a very expensive hardware in order to really make use of Datacenter. Most companies with a decent size number of clients are probably going to be purchasing Standard.

Minimum Requirements

4:25-5:18

The minimum requirements for all four versions are very simple. Your CPU has to be at least 1.4 GHz. It's got to be a 64-bit CPU. They dropped support for 32-bit server operating systems in server 2008 R2. I really like this; your minimum RAM is 512 MB. Now, in reality, you're probably not going to install a server with just 512 MB RAM. This makes it great for testing. If you want to go through and build a test server, play around with some of the technology, you actually don't need a huge amount of RAM. You could even get a little netbook with a gig and a half of RAM, and run one virtual machine, and play with server. You need at least 32 GB of free space on the hard drive for the system volume. That's for the operating system alone. Certainly, you want to leave a lot more space than that, because over the course of time, there's going to be patches, there's going to be upgrades, and then it's probably going to do more than just hang out and be a server.

Upgrade Paths

5:19-6:21

If you have an existing Windows Server operating system, there are certain upgrade paths. There used to be different editions of Windows Server 2008 or 2008 R2. If you have Standard or Enterprise with SP2, you can upgrade to 2012 Standard or Datacenter. There's no more Enterprise version with 2012. If you've got Datacenter, you've got to go up to Datacenter. There used to be a web version that only did web stuff. That is not being supported anymore with 2012, so you would go to 2012 Standard. Then, pretty much the same thing with Server 2008 R2. A good rule of thumb with upgrades is, you can go to your edition, or better when you're upgrading. That's kind of a generic Microsoft rule. The only exception I see here is that, technically, it looks like it's possible to go from Enterprise to Standard. I can't image anybody that bought Enterprise is going to go to Standard. You're probably going to go to Datacenter. If you're looking for something easy to memorize for a test for the rest of your Microsoft career, generically, you can go to an equivalent edition or better, usually from the most recent operating system.

Test Tips

6:22-6:47

Make sure for the test, you notice, this is 2008 with Service Pack 2. If it's 2008 R2, it's Service Pack 1. In Microsoft tests, sometimes they try to be tricky. They'll say, 'Oh, you have Windows Server 2008.' What they might be looking for you to say in the answer is, well, I need to upgrade that to SP2 in order to do the upgrade. Just make sure that you're aware of those things, so that you can't just, well, they tossed on an SP there, that might be relevant for a test question.

Installation Options

6:48-7:27

Finally, when you go through the install, the last thing you're sort of going to choose is to do a Full Server install with a Graphical User Interface. OK, GUI means Graphical User Interface, or just a Server Core install. We are going to cover those in separate modules so you understand exactly what that is. The really cool thing about Server 2012 is, we can switch between them.

If you feel like you can't make a decision, no worries, you can change your mind after the fact. That is the information that you need to keep in mind when you're installing Server 2012, or if you're upgrading. Make sure you have an understanding of your upgrade paths and you know your minimum requirements. Once we got the server installed, we'll be able to jump in and play with some of the cool stuff that comes with it.

2.1.2 Installing Windows Server 2012

Installing Windows Server 2012
0:00-0:03
This is a demonstration on how to install Windows Server 2012.
Windows Setup Screen
0:04-0:16
We have a hard drive with nothing on it. Put the DVD in, and boot up. The first thing it's going to do is come into the Windows Setup screen. We'll go ahead and click Next. The next thing we're going to click is Install Now.
Repair Your Computer Link
0:17-0:30
I do want to call your attention to this Repair your computer link in the initial setup. If you do have a failed server, and you need to restore from a backup, this is the way you would boot into the Windows recovery environment.
Install Now
0:31-0:35
We're just going to click Install Now.
The first thing we need to do is put in our product key.
Product Key
0:36-0:49
Whichever product key you put in is going to govern which version of Windows gets installed. Once you get your product key in, you click Next.
Installation Options
0:50-1:28
Now, we have our installation options. We can go through and install this with a Server Core installation, which is highlighted by default. That's just going to give me the command prompt with no GUI, where we're actually going to go ahead and install a server with a GUI. Click Next. Of course you have to "Accept the License Terms".
This is a neat screen because both Upgrade and Custom Install are activated, but this is a blank hard drive. There's nothing to upgrade. If I click this, nothing is going to happen. This would be used for...if I'm going to run the DVD and upgrade an earlier version of Windows. Something is already running on this computer.
Custom: Install Windows Only (Advanced)
1:29-1:36
For a clean install, you always go to Custom: Install Windows only (advanced). We're going to click on that. We choose where to install it.
Choose Where to Install It
1:37-1:44
This particular computer just has one hard drive, so we're going to install it on the Unallocated Space.
Load Driver
1:45-2:02
If for some reason your hard drive does not show up in here, it's usually because there's no driver. You would click the Load Driver link and load the driver. If it's some kind of SATA or RAID card, put the driver on a USB or CD and click Load Driver, and then you'll be able to see the hard drive.
Partition the Hard Drive
2:03-2:30
If you did want to partition the hard drive, we can go to Drive options (advanced), and we have the option to make a New partition. If any partitions already existed, we could Delete them, Format them, Extend them; whatever we need to do to work with the hard drive. I just want to install it on the unallocated space, so I'm going to click Next. It's got several things it goes through. It tells you, right here, it might restart a couple of times. Then, once it's done that, we will be ready to set the administrative password and get into Windows.
Administrator Password
2:31-2:54
Once it's gone through the setup, you really don't have to be involved with it after you've identified where on the hard drive to install it, until it comes up and it asks you for the administrator password. By default, this should be a complex password. You don't type something it doesn't like, it'd certainly tell you. Once you've got your password typed in, click Finish. It's going to finalize your settings. The next thing that we should see is Server Manager.

Sign In

2:55-3:07

Actually, we're going to hit Control, Alt, Delete, in order to sign in. There's Server Manager, which opens up automatically when we log in for the first time.

Local Server

3:08-3:51

What you should do after installation is make sure you come in to Local Server, and because it didn't ask us any questions during installation, it's going to have a randomly generated computer name. Certainly, it's part of a WORKGROUP. The Time zone is usually set to Pacific Time, because that's where Microsoft is. If nothing else, you should at least make sure you configure the Computer name and the Time zone. You need to set up the network settings. You can do that from in here.

The big one people forget to do is the Computer name, because once you join the domain, it becomes a little more difficult to change the name. It's best to do that right off the bat. Get the Computer name and the Time done. Then, you're in good shape. That's how you install Windows Server 2012.

2.1.3 Server Installation Facts

Windows Server 2012 is the latest release of the Windows Server product. With a Metro user interface similar to that of Windows 8, the installation of Windows Server 2012 differs from previous versions of Windows Servers in the following ways:

- Editions of Windows Server 2012 support only 64-bit processors.
- You must enter the product key when you are installing the product.
- There are four editions of Windows Server 2012.
- All editions are on the same installation DVD. The product key determines which edition is installed.

When choosing an edition of Windows Server 2012, make sure that you are aware of the features that it supports and select the edition that meets the needs of your organization. The following table lists the Windows Server 2012 editions.

Edition	Features
Foundation	The Foundation edition is a scaled down version of the server operating system, used for small businesses and supports most server roles. Foundation supports: • One CPU socket and up to 32 GB of RAM • Up to 15 simultaneous LAN connections • Up to 50 Routing and Remote Access (RRAS) connections • Up to 20 simultaneous Remote Desktop (RD) connections Foundation does *not* support: • Server Core installation • Hyper-V virtualization services • Server clustering
Essentials Edition	The Essentials edition provides additional hardware support and role support above what is provided by the Foundation edition. Essentials supports: • Two CPU sockets and up to 64 GB of RAM • Up to 25 simultaneous LAN connections • Up to 250 Routing and Remote Access (RRAS) connections • Up to 250 simultaneous Remote Desktop (RD) connections Essentials does *not* support: • Server Core installation. • Full Hyper-V virtualization services. The Hyper-V host can be installed into Hyper-V as a guest. • Server clustering.

Standard Edition	The Standard edition is available for medium and large businesses. Standard edition supports: • Up to 64 CPU sockets and 4 TB of RAM • Unlimited LAN connections • Unlimited simultaneous Routing and Remote Access (RRAS) connections • Unlimited simultaneous Remote Desktop (RD) connections • Server Core installation • Hyper-V: two Hyper-V guest OS sessions using the same license key as the host (provided the host is *only* running Hyper-V) • Server clustering with up to 64 nodes per cluster
Datacenter Edition	The Datacenter edition provides much more capability than the other Windows Server 2012 editions. • Up to 64 CPU sockets and 4 TB of RAM • Unlimited LAN connections • Unlimited simultaneous Routing and Remote Access (RRAS) connections • Unlimited simultaneous Remote Desktop (RD) connections • Server Core installation • Hyper-V: Unlimited Hyper-V guest OS sessions using the same license key as the host • Server clustering with up to 64 nodes per cluster

The following table identifies the minimum requirements for all editions of Windows Server 2012.

Component	Minimum
CPU	1.4GHz (x64)
RAM	512 MB
Free space for System Volume	32 GB

The following table identifies Windows Server 2012 Upgrade paths.

If you are running these editions	You can upgrade to these editions

Windows Server 2008 Standard with SP2 or Windows Server 2008 Enterprise with SP2	Windows Server 2012 Standard or Windows Server 2012 Datacenter
Windows Server 2008 Datacenter with SP2	Windows Server 2012 Datacenter
Windows Web Server 2008	Windows Server 2012 Standard
Windows Server 2008 R2 Standard with SP1 or Windows Server 2008 R2 Enterprise with SP1	Windows Server 2012 Standard or Windows Server 2012 Datacenter
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2012 Datacenter
Windows Web Server 2008 R2	Windows Server 2012 Standard

In addition to the various server editions, you can install Windows Server 2012 as a Server Core installation. Server Core is a minimal server installation option which does not provide a graphical user interface. Unlike previous versions of Windows Server, you can switch between a Server Core installation and a full installation.

Once the software has installed, Server Manager opens when you log in for the first time. Using Server Manager:

- Provide a computer name
- Set the time zone
- Configure network settings

2.2 Powershell

As you study this section, answer the following questions:

- What is the typical PowerShell command format?
- How can you access the cmdlet help system?
- How can you cycle through options for a cmdlet?
- How can you use the auto-complete feature of PowerShell 3.0?
- What is the purpose of Windows PowerShell providers?
- How can you import a PowerShell module when writing a script?

After finishing this section, you should be able to complete the following task:

- Use PowerShell cmdlets to configure a Windows Server 2012.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  Configure Server Services

This section covers the following 70-410 exam objective:

- 102 Configure servers.
  Install and configure Windows PowerShell Desired State Configuration (DSC)

2.2.1 Powershell Overview

PowerShell Overview
0:00-0:02
We'll go through an overview of PowerShell.

Move Toward Command Line and Scripting
0:03-0:43
There's been a trend in the past, I would say, five years--moving more towards command line and scripting. As networks get bigger and bigger, it becomes more time economic to create scripts or batch files that do routine tasks over and over again. It's fine to work with the Graphical User Interface if I need to make a change to one user, but what happens when I need to make a change to 150 users, or I need to make that change remotely on 10 particular workstations? In that case, the command line in PowerShell becomes much more powerful. We're going to talk a little bit about some of the things you need to know about PowerShell. I'll even toss in a little bit of information about the command line, because that also can be pretty useful.

Command Line Interface
0:44-1:56
The command line interface -- a lot of students will say to me, "Shad, how do I know how to do the command line"? What I love about the command line is, it works very similar to English. You have a verb, and then you have an object--you might think, why are we getting an English lesson? Believe me, this is some powerful information. And then, if anything at all, you have some adverbs out at the end. The thing to know is, there's a space between each piece. If I were doing a command, for example, where I'm going to use a Dsadd command to add a User to Active Directory, and then my adverb modifying the verb might give information like, what user I'm going to add--you know, cn=shad and so on and so forth.
The thing to realize with the command line is, we're always starting out with a verb. There's a space, there's some type of object, what the verb is going to act on, and then anything else that comes in is going to be an adverb. Sometimes, switches have a slash (/) in front of them. Sometimes, they have a dash (-) in front of them. I might do a show me a Directory Listing of a folder with a /w for wide. Some commands you have to use a slash, some you have to use the dash, some don't care whichever one you use. Once you get familiar with commands, you'll get an idea which one you actually need to have.

PowerShell
1:57-3:28
PowerShell is like a command line on steroids. The command line really comes out of the old DOS years. They have beefed it up a little bit for Server, but still, and for all, there's not that many things that you can do from the command line. There's been a movement in the last three generations of Windows Server, to have the whole thing based on PowerShell. Even when you're using a GUI interface like the DNS snap-in or the DHCP snap-in, or particularly in the other Microsoft products like Exchange or SharePoint, you might go in Exchange and create a mailbox -- what it's actually doing is running a PowerShell command in the background. In some of the interfaces, they'll have a button where you can see the PowerShell command. Others, they haven't quite caught up yet. My assumption is, they're going to be doing that in the future.
Even the GUI snap-ins that you're using are working on PowerShell in the background. The rule is that there's really nothing that you can do in Server 2012 that you can't do from PowerShell. PowerShell could run the entire server, without any GUI at all, which will get us to Server Core, where we won't even have a GUI. The only difficulty then is learning the PowerShell commands. The other thing I'll say about PowerShell is, there are some things that can only be done in PowerShell, or that can actually be done more easily in PowerShell. If you get an idea of how PowerShell works generically, which is what we're going to do in this lesson, as you go through the course and we talk about specific PowerShell commands, you'll have a better idea of understanding how those commands will work.

PowerShell Syntax
3:29-4:55
PowerShell commands always start out with a Verb-Noun. For example, I could generically go in there and say, "Hey, I need to Get-help", or "I need to make a New-ADuser", then there'll be a space, and then come in, your adverbs, and that type of thing. Generally, the commands themselves are not case sensitive. The switches themselves are not case sensitive. I'm not sure if this is the right syntax, but let's say a switch here is -Name and then you do a space and you put the name, whatever you're going to do. Usually, the switches are not case sensitive.
If you're turning on a feature-- let's say, I'm going to use PowerShell to turn on DNS--that type of thing would be case sensitive. If it makes you uncomfortable, then just stick with the case that they give you in the Help system; if not, you can go through and mix up the case and see what works. Generically, New, /Add, and /install are always for working with a new object. Get is always to retrieve information about an existing object. This is something that's

already there, and you want to know about it. Then, set is to make a change, again, to an existing object. Besides those basic principles: /Add, New, /Install, Get, Set, a couple of also overriding principles with PowerShell.

The "Sarcastic Answer"

4:56-5:36

First of all, whenever you're looking at something Microsoft, I always tell my students, if you have to guess in an exam, guess the sarcastic answer, so if somebody says, 'Well, what would be the PowerShell command to make a new Active Directory user?' You want to be a little sarcastic and say, 'Well, duh, it's probably New-ADUser.' Yup that's the right answer. You can always guess the sarcastic answer. A lot of students will talk themselves out of a correct answer, because they think it can't be that easy. If it was that easy, why would it be on an exam? They try to make the commands easy, so that you can figure out what they are without having to have extensive training in PowerShell.

If you're looking for a command that does something, go ahead and try the one that's the most obvious.

Finding All Available Options

5:37-6:03

The second general principle is, if you type New, let's say, and a dash (-), in PowerShell you can hit the Tab key and it will scroll through all the available options. If you're wondering, can I make a New ADUser? You can type New, type a dash (-), and then hit Tab to scroll through. If it's in there, you can do it. If it's not in there, then you can't do it, and try a few different verbs to see what's in there. So, PowerShell itself will help you to understand it.

Flavors of PowerShell

6:04-7:36

The last thing you want to be aware of is, there's not just PowerShell. There are different flavors of PowerShell, believe it or not.

The generic PowerShell that you might launch from the taskbar is just that--a generic PowerShell. As you add functionality into Windows Server, there will be new PowerShell commands that become available. Those are in modules. So, let's say I install Active Directory. There's actually going to be added in, an Active Directory module for PowerShell. There'll be a separate icon for that module PowerShell in the Start menu that I can click on. In that PowerShell will be all of the Active Directory commands. If I open the generic PowerShell, those commands won't be in there. I have to go into the module-specific PowerShell.

This becomes a little bit of an issue if I want to save a PowerShell script, so if you're in the generic version of PowerShell, you've installed something that has added functionality, then you can use the command import-module, and then specify the module that will be imported. That's exactly what you would do at the top of the script. If I wanted to save a script into a text file, so that I can run it anytime I want to run it, at the very top, I would say import-module whatever the name of the Active Directory module is, and then I can go ahead and use those Active Directory commands, then there would be no problem, they would have been imported into the generic PowerShell. Just be aware of that as well, that there are different modules that will extend the functionality of PowerShell, so that really anything Microsoft that gets installed into the operating system ultimately can be configured using PowerShell.

2.2.2 Using Powershell

Using PowerShell
0:00-0:03
This is going to be a demonstration of using PowerShell.

Getting into PowerShell
0:04-0:19
The easiest way to get into PowerShell is to click the icon on the Task Bar. We're going to go down into the lower left hand corner of the screen and just click the PowerShell icon. It should open us up into the generic PowerShell window.

PowerShell Commands
0:20-1:05
PowerShell commands are always a verb, a dash, and then a noun. If you don't know what the exact command is that you need, you could always try typing a verb, and then hit Tab and it will cycle through all the different things that you can do that start with that verb. New-, Add-, and Install- are always for working with new objects ... Active Directory domain, Active Directory domain controller. Get- is always for getting information about existing objects, and Set- is always for changing the properties or settings of an existing object. If you don't understand what a command does, you can always use the get-help and then put the command that you're not sure what it does.

get-help
1:06-2:38
I could do get-help on a new-aduser. PowerShell commands are not case- sensitive unless you're using it to install a role or something like that, in which case, the name of the role or feature would be case sensitive, but not the PowerShell command itself.
It gives me quite a lot of information; up at the top it's got the entire syntax, which I don't usually find helpful. But here's my New-ADUser. These are all the different things I can specify about the user. The great thing is, here's an example of a command; so they're giving you an example of making a new user the SamAccountName, which is the login name of "glenjohn", -GivenName which is a first name "Glen", -Surname of "John", -DisplayName of "Glen John". If there's any spaces, you've got to put that in quotes. Path tells it where in Active Directory to create it, so the Microsoft help is saying in the Users container, in the fabrikam domain, in the local domain, and then if you have any other attributes, you can go ahead and set those up there too. You must at least specify the SAMAccountName parameter to create a user; so at the very least they need a logon name and it will go through. Then it tells you if you want more examples, you can do your get-help new-aduser and add a -examples, and they give you more examples of how you would do that.

More Examples
2:39-2:53
You literally could come in and go into the upper left hand corner, click Edit, choose Mark, and now I can highlight one of these that I like.

Using the Clipboard
2:54-3:40
Let's say, my command should be pretty similar to this one.
Once I've got it highlighted I hit Enter on the keyboard that puts it into the clipboard; so I can go back up the upper left hand corner, Edit, Paste it in, and it will put that command up, and then I can edit it as much as I want. So I can either type it from the screen from the examples that it's giving me, or I can copy and paste it right back into PowerShell, and then edit it as much I need to.

Commands Specific to Installed Applications
3:41-4:57
If you open up the generic PowerShell and you don't have commands available to you because you're looking for commands that are specific to an application you've installed like Active Directory, Exchange, SharePoint; when you've installed an application that has specific PowerShell modules in the Tools menu in Server Manager or in the Administrative Tools menu off the Start menu, you will sometimes see a particular module. This Active Directory module for PowerShell opens up a PowerShell window with the Active Directory module imported by default, and if I install Exchange, I'd see one for Exchange; if I install one for SharePoint, I'd see one for SharePoint. If you're ever missing any commands that would tell you that you don't have the module running and you can go into Tools, you can launch the specific PowerShell module for that application or in the generic PowerShell, you can use an import-module, and then you would have to specify the particular module you want to import. If you're going to make a script with PowerShell that's dependent on particular commands that are in some of those modules, you have to add that import module with the correct module name to the top of the script.

2.2.3 Powershell Desired State Configuration (DSC)

PowerShell Desired State Configuration (DSC)

0:00-0:21

In this lesson we're going to spend some time looking at the PowerShell Desired State Configuration feature or DSC. Basically, what DSC does is provide you with some advanced scripting capabilities for your Windows Server system. DSC is a PowerShell extension and it ships with Windows Server 2012 R2 as well as Windows 8.

DSC Usually Pushes Changes to the Target Systems

0:22-1:27

DSC is really cool. It makes the life of a system administrator much easier. One of the key reasons why is the fact that DSC can use either push or pull operations.

Really DSC can do a lot of the same things that Group Policy already can. The key difference however is the fact that Group Policy always uses a pull operation, meaning that the target system has to initiate the configuration process. For example, if you are working with a computer policy, when does that computer policy get applied? It's whenever the system powers on. Likewise with a user policy. When does it get applied? When the user logs in. If you make a Group Policy change after a system is booted or after the user has logged in and you want to apply it, you have to run GPupdate on the system to pull that information down from the domain controller.

Using DSC we can change things around a little bit. We can actually push changes out to the target systems. DSC can also use a pull model if you want, but most of the time the push model employed instead.

DSC Functions

1:28-2:16

Let's take a look on how DSC works in a little more detail. Now DSC can do a lot of things. For example, it can install or remove server roles and features. You can use it to manage registry settings. You can use it to manage files and directories in the file system. You can stop, start, and manage processes and services on the server. You can manage local groups and user accounts. You can install and manage packages, including both .MSI packages as well as .EXE files. We can manage environment variables. We can run PowerShell scripts. We can fix a configuration that has drifted away from the desired state. And we can discover the actual configuration state of a particular host on the network.

DSC Installation

2:17-2:35

Be aware that DSC must be installed first before you can use it. It's a Windows feature. You use Server Manager to install the DSC feature. Also being aware that DSC uses WinRM. The WinRM service must be enabled on all the target systems that you're going to manage with DSC.

DSC Providers

2:36-3:49

DSC uses providers. These providers basically make things happen on the target systems that you're managing with DSC. It can a lot of things.

As you can see here we have the archive resource, which can be used to unzip files. We have the file resource, which can be used to copy files or even as noted over here, you can copy an entire tree structure if you need to. We have the group resource provider, which manages our local group memberships. We have the user resource, which you can use to create modified local user accounts. The package resource, which is used to manage software packages. The process resource, which is used to actually run applications. We have the registry resource, which can be used to modify the registry. We also have the role resource, which can be used to install a Windows role or feature. The script resource, which can be used to execute a PowerShell script. We have the service resource, which we can use to manage services on the system. We can stop a service. We can start a service. We can set the startup type and we can set the startup account as well.

Finally, we have the log resource, which is used to send events to the Windows event log for troubleshooting purposes.

Creating a DSC Configuration

3:50-4:05

With that in mind let's spend a little bit of time looking at how to create a DSC configuration. In order to do this we're going to look at a simple DSC script that's going to install the IIS Web service on a target system.

Declare the DSC Configuration

4:06-5:13

Now in order for a DSC configuration to be consumed, you first have to declare it within your script file. An example of doing that is shown right here.

We simply add the text that you see here: Configuration AddIIS. Now this word right here is simply a name that I defined. It's not a keyword. You can define whatever you want. I chose AddIIS because that's what we're going to

do with this configuration. We're going to tell the target system via DSC to install the IIS role. Also notice that the PowerShell DSC extension adds a new keyword. It's called 'Configuration.' You must include this in order to declare your DSC configuration here. Now this configuration keyword really acts more like a function and if you've done any scripting at all, you say "It looks like you're declaring a function." Well yeah, that's really what you're doing with the Configuration keyword with DSC.
We have configuration, the name of the DSC configuration that we are declaring and then we have our opening and closing curly bracket.

Extend the Configuration

5:14-5:51
Once we've declared our DSC configuration, we can then go ahead and extend the configuration. In this example, what we're going to do is add the IIS role to the target system. Notice the first thing we do is define a parameter called $MachineName. What we'll do is when we run this script, we'll supply the name of the host that we want this configuration to be applied to and the name of that host will be put into the $MachineName variable.
The next line says that everything that comes after this is going to be performed on the host that is contained in the $MachineName variable.

Comment

5:52-5:54
Then we have a comment that says we're going to install the IIS Role.

Specify the DSC Provider

5:55-6:41
Next I need to specify which DSC provider that I want to use. In this case I want to use the WindowsFeature DSC provider. Remember as we talked about just a second ago, WindowsFeature is used to install a role or feature on the target system. In this case we're going to install the IIS role on the target host, whichever one we specify here using the WindowsFeature DSC provider. Note down here we have the Ensure parameter set to a value of "Present" meaning that we want this role to be present on the target host. So if it is, then we won't actually install the role, but if it is not then the WindowsFeature provider will go ahead and install the IIS role on the target system.

Create a DSC Consumable MOF File

6:42-7:36
Now that we have declared the DSC function, we next need to create a DSC consumable file. It's called an MOF file. We do that by calling the function that we just created. We do that just like we would any other PowerShell function. In this example remember we declared a DSC configuration named AddIIS. We're going to call it and then we're going to specify a MachineName that's going to be fed into the $MachineName variable that we looked at on the previous slide and the name of the host that we want to use is WEB1. This will cause the lines that we configured on the previous page to be pushed down to the network host with a name of WEB1.
Now by doing this, a folder is going to be created. The name of the folder is going to be the same name as the DSC configuration. In this folder the MOF output file is going to be created.

Applying a DSC Configuration

7:37-8:32
At this point we have created our MOF file. That MOF file can be consumed by DSC. We can apply it using a command similar to what we see here. We run the 'Start-DscConfiguration' PowerShell commandlet and then we specify the '-Path' parameter and then we specify the name of the folder that was created on the previous slide. Remember we created the AddIIS folder. This folder contains the MOF file that DSC can use, can consume if you will, and we added the '-Wait' and '-Verbose' parameters as well.
By running this command DSC will locate the host that we specified earlier, WEB1, and then it will use the WindowsFeature DSC provider to see whether or not the IIS role has been installed on that host. If it has, then it won't install the role obviously, but if it has not, then DSC will go ahead and install the IIS role on that host. As you can see DSC is pretty useful.

Summary

8:33-8:57
I didn't have to wait for Group Policy to apply the configuration change to the system. Nor did I have to go out and visit the console of the WEB1 system in order to make the change. That's it for this lesson. In this lesson we talked about DSC. We talked about what DSC is and what it does. We looked at how it works and then we made a sample DSC configuration that would dynamically install the IIS role on a target system on our network.

2.2.4 Using Powershell Desired State Configuration

Using PowerShell Desired State Configuration (DSC)
0:00-0:32
Let's take a look at using DSC to make some configuration changes. This is going to be the server on which I run DSC. Essentially my script is going to do two things. It's going to take a file inside the DSC folder which I've shared called test.txt and copy it to the target server. It's also going to go ahead and install IIS on the target server. I'm on DC1 where I'm going to be creating and running my script. Let's take a quick look at our target server.

Target Server
0:33-0:53
My target server is named DirSync and if we look, it has a local folder named DSC which I've also shared out. Currently there's nothing in this folder. If we see the test file show up, we'll know that that piece of the script has worked.

Setting Up the DSC Server
0:54-0:59
Now let's go back and set up our DSC server. I'm back on DC1.

Install DSC
1:00-1:33
The first thing I want to do is install DSC. I'm going to Add Roles and Features and hit Next until I get to my Features. If we scroll down, there's a PowerShell section and if I expand that, you can see that I can install the Windows PowerShell Desired State Configuration. It's going to pull in a whole bunch of features that are required, say 'yup', Next, Next, Next, and then Install. Once I have installed the feature, my first step is going to be to create my configuration file.

Create the Configuration File
1:34-3:46
I've already typed this out. I use the keyword 'Configuration' and then I put whatever I'm going to call my configuration. 'Node DirSync' specifies the server that this is going to take place on. You can add comments using the pound or the hash sign. This section here is going to copy a file, so the keyword really is 'File' and then I can call it anything I want. I've got a set of brackets. Make sure that you use the braces, not square brackets. Ensure = "Present" is going to turn it on. If I wanted to delete it I could change "Present" to "Absent". I've got the share that the file is coming from which is on DC1 and then my destination path is where it's going to. This closing bracket ends that section. Then I've got a section that's going to install IIS on the destination server.
My resource is WindowsFeature and it's installing IIS. This is why you need to look up the resources. It'll tell you the different parameters that are required. In this case I need Ensure = "Present" or "Absent" which will install it or remove it. In this case I'm going to be installing it so I've got Ensure = "Present". Then I need to give the name of the feature as I would get it from PowerShell. If you do a get Windows Feature, it's actually named Web-Server. This bracket closes the piece for IIS. This bracket is going to go ahead and close all the different resources that I'm using and then this closes the entire configuration. Then this last DSCTest makes a sub-folder where the MOF file is going live. I'm going to go ahead and save this and I'm going to save it in the C: drive in this DSC folder and I'm going to name it 'DSCTest.ps1'. I put the quotes so it doesn't get an extra .txt. If I go in here, I can see my configuration file.

Create the MOF File
3:47-4:34
My next step is to go into PowerShell and create the MOF. We've created our MOF file in the \dsc\DSCTest folder. If I go in here I can see that folder's there with the MOF inside of it. One more thing I want to check on my target server--just so you can see it actually works. We'll just hop back to DirSync for a moment. We're over here on DirSync and I'm just going to quickly go in and click 'Add Roles and Features' and hit Next until we get to Roles. You can see that right now the Web server is not installed, so that'll be something else that we can check. We'll check both that and we'll be checking in the C: drive in DSC to see that our file gets copied over.

Execute the Configuration
4:35-5:16
Let's go back and execute our configuration. Go back into PowerShell and I'm going to use the Start-DscConfiguration -computername is the name of the target, in this case DirSync. The path is the path to the MOF file. Since I'm in the DSC folder, it's going to be DSCTest. I'm going to tell it to wait until everything is done and -verbose will give me the maximum amount of information. Let's take a look at what happens. It's copying the file and you can see at the top 'Start Installation' and that's going to go until 100% of the Web server has been installed.

View the Results on the Target Server

5:17-5:49

Now it's done, we can see the output, but really the best thing is to go back to our target server and take a lot at what happens.

I'm back on DirSync. Nothing really happened in the console. But, if I go up to the dashboard and I start to Add Roles and Features and I Next, Next, Next, until I get to the Roles. We can see that in fact the Web server is now installed. If I go into my DSC folder, there's the Test document that got copied over.

Summary

5:50-6:07

DSC can be used for a whole lot of different things. Many of those things can be done using Group Policy, but one advantage of DSC is that you can push configuration changes over and manually store the execution of that, as opposed to having to wait until Group Policy gets reapplied.

2.2.5 Powershell Facts

PowerShell is a powerful scripting tool that you can use to automate system administration and application management tasks from the command-line. Using PowerShell, you can create scripts to manage the Windows operating system and applications that run on Windows. PowerShell commands enable you to perform certain tasks, especially those involving a number of objects, much faster than the graphical user interface (GUI). Some tasks can only be done through PowerShell.

The following terms are used with PowerShell:

Term	Definition
Command line	The *command line* refers to the command-line interface (CLI) used to interact with a computer program. Commands are entered at the command prompt. The command-line interface: • Is text based. • Uses a simple command construction, also referred to as the command *syntax*. • Uses *flags*, also referred to as *switches*, to specify an option in the command.
Cmdlets	*Cmdlets* are commands that an administrator enters at the PowerShell prompt to perform system management tasks. The PowerShell command-line interface uses simple command construction, typically in the form of: **[verb] -[noun] -[adverb]** When using PowerShell cmdlets, keep in mind the following: • The cmdlets are generally not case sensitive. • The adverbs are typically referred to as switches and are not case sensitive. • If you are using PowerShell to turn on features, use the same case used in the cmdlet help system. • Common cmdlets include: **new**, **add**, and **install** are used when working with a new object. **get** is used to retrieve information about an existing object. **set** is used to make a change to an existing object. **get-help [cmdletname]** is used to get help for a particular cmdlet. • PowerShell cmdlets are designed to be lightweight, easy to recognize, and easy to use. • If you type a PowerShell verb followed by a hyphen (for example **new -**) and press the TAB key, you can cycle through the available options.

	PowerShell 3.0 has IntelliSense capabilities that allow you to auto-complete commands by pressing the TAB key. • The PowerShell icon on the taskbar opens the generic version of PowerShell. • Installing services and features on Windows Server creates new versions of PowerShell that contain modules with the cmdlets for the added services and features. There is a separate icon to launch the new instance of PowerShell. When you write a script that requires the cmdlets from an added module, you can add the cmdlets by entering: **Import-Module [modulename]**
Providers	Windows PowerShell *providers* allow you to access data stores, such as the registry and certificate store, in a way similar to accessing the file system.

Some new features in Windows PowerShell 3.0 included with Windows Server 2012 are:

- The Deployment Image Servicing and Management (DISM) module **Get-WindowsOptionalFeature -Online** cmdlet allows an administrator to view features installed on a remote computer. Administrators can then add or remove cmdlets from the remote computer using the **Add-WindowsFeature** and **Remove-WindowsFeature** cmdlets.
- Remote Windows PowerShell connections have been improved to attempt to maintain a connection even when network connectivity is lost. Previous versions of PowerShell easily dropped the remote connection when network connectivity problems were encountered. PowerShell 3.0 allows you to reconnect from a different computer if the connection is lost.
- Windows PowerShell Workflows are sequences of multicomputer management activities. Windows PowerShell Workflows allow workflows written in XAML or the PowerShell scripting language to be run as a cmdlet.

Windows Server 2012 R2 introduces PowerShell Desired State Configuration (DSC) which allows administrators to manage the configuration and environment of computers and devices connected through a cloud infrastructure. With DSC, the administrator uses declarative scripting to specify and manage configuration. DSC resources include the following:

- Management Object Format (MOF) files are used to specify management tasks to be performed. Administrators can create MOF files in a variety of ways, including PowerShell v4 declarative syntax extensions or third-party tools. PowerShell v4 has built-in resources to facilitate MOF creation. These resources include:
  WindowsFeature identifies the role, such as Web server.
  File identifies and manages files and directories.
  Group identifies and manages local Windows groups.
- PowerShell keyword configuration is used to create MOF files.
- The creation of custom resources using PowerShell v4 or third-party tools is supported.
- DSC Local Configuration Manager runs on all target nodes and calls the necessary configuration resources.
- The option to use a push or pull implementation.

In the Pull Model, DSC data and custom providers are kept on an IIS Web server. The target contacts the IIS Web server to obtain the configuration instructions.
In the Push Model, the configuration instructions and custom providers are pushed to the target system.

2.3 Server Roles

As you study this section, answer the following questions:

- What is the relationship between server roles, role services, and features?
- How is a new role installed on a server?
- How do you remove a role in Windows Server 2012?
- Which role should be installed if you want to configure a local update server?
- How can you migrate roles from an existing system running an earlier version of Windows Server to Windows Server 2012?
- Which Windows Server 2012 feature allows you to remove the source files of unused roles and features?
- What is the benefit of removing the source files of unused roles and features?
- What methods can be used to obtain source files to re-install removed roles and features files?

After finishing this section, you should be able to complete the following tasks:

- Install roles on Windows Server 2012.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.

This section covers the following 70-410 exam objectives:

- 101 Install servers.
  This objective may include but is not limited to:
  Plan for server roles
  Optimize resources utilization by using Features on Demand
  Migrate roles from previous versions of Windows Server
- 102 Configure servers.
  This objective may include but is not limited to:
  Configure services

2.3.1 Server Roles

Server Roles
0:00-0:35
Let's talk about Windows Server 2012 Server Roles. Server roles are a concept that came in with Windows Servers 2008. Basically, what happened was Microsoft looked at how servers operate, and they found out that in companies where you have different servers doing different things, they take on a role, like maybe this server is a DNS Server or this server is a DHCP Server.
Instead of having everything be a feature that gets added or removed, they set up roles which are major functions of the server: like DNS Server, DHCP Server, Active Directory Domain Controller.

Roles and Role Services
0:36-0:46
When you add or remove a role, often there will be subsets of the role which are called Role Services. It doesn't really matter memorizing whether something's a Role or a Role Service.

A Single Wizard
0:47-1:52
The really cool thing about Windows Server 2012 is, it's all one wizard now. In Server 2008 and 2008 R2, you had roles, and then you might have to go in and add Role Services using a different part of the GUI.
There was also something called Features, which are still in Windows Server 2012. These are not major functions of the box. For example, clustering is not a major function of the server. You would be clustering something like DNS, DHCP, file services, so that's added in as a Feature.
Features are in the same wizard as roles. It's really a one-stop shop to go in and add either a Role or a Feature to the server. The thing that is a little bit weird with Windows Server 2012 is, if you have to remove a role or a feature, then you have to launch the Add Role or Feature wizard, and then there'll be a hyperlink really early in the wizard, that says, "If you're here to remove a role click this hyperlink", which will then launch the Remove Role or Feature wizard.
That's the only part you have to watch for. There's not an actual command, 'Remove Role or Feature'. You can add a role or feature and then remove it, but we'll take a look at that in the GUI.

Summary
1:53-1:24

Removing a Role or Feature
1:25-2:02
Role is a major function of the box, Role Services subsets of roles, and then Features supporting the roles of the server, all added in using the same wizard.

2.3.2 Migrating Server Roles

Migrating Server Roles
0:00-0:44
In this video, we're going to talk about migrating server roles. We're going to take a generic look at it. It's very specific to the role. If you really do have to migrate a server role, you're going to have to dig in a little bit deeper than we can go, because we don't really have any particular roles we want to pull out, to focus on.
There are some tools inside of Windows to help you with migrating. In most cases, rather than using the migration tools, it would be easier to upgrade the existing server or have a new server take over, run them side-by-side and then decommission the old server.
Certainly, I want to make you aware of the fact that tools are out there, because it's possible they may prove useful to you.

Installing Migration Tools
0:45-1:04
We're going to start by installing the migration tools. We can either do this using Add Roles and Features -- and I will show you that -- Windows Server Migration Tools, or it will be more fun to use PowerShell, so we're going to do that.

install-windowsfeature migration Command
1:05-1:47
We do an install-windowsfeature. The feature we installed is migration. It goes through; turns on the migration tools. All the warning is saying is that we don't have Windows Update turned on. If there are updates for the migration tools, we're not going to be receiving them, but the feature was installed successfully.
Our migration tools will show up on the Start menu. You see that right here -- Windows Server Migration Tools. What it does is open up a special PowerShell prompt that has those tools already loaded in.
Now that I've installed the migration tools on the server that's going to be receiving the role, we're in 2012, we're going to be receiving it.

Creating a Package
1:48-2:33
I need to create a package that's going to be used to deploy these same tools on the computer that has the role right now, the one I'm migrating. We're going to do that from the command prompt.
The first thing I want to do is change directory (cd) to where the migration tools are kept. These are located in the c:\windows\system32\servermigrationtools directory. To create my package, we're going to use a command called SmigDeploy. As a general rule, if it's something to do with migration, it usually has a mig somewhere in the command.

32-bit or 64-bit
2:34-3:24
Architecture is what type of operating system it is, either 32-bit or 64-bit. If I'm going to be deploying this to a 32-bit version of server, I'm going to say X86, we're going to be using a 64-bit, and the code for that is amd64. That's for any 64-bit processor. It doesn't have to be AMD, it just means it's 64-bit.
The operating system running on the server that we're going to be using to migrate is Windows Server 2008 R2. You've got to put the appropriate code in, so if it was Windows Server 2008, I would leave off the R2. If it was actually 2003, it would be WSO3. We're going to put it in a folder in the C: drive named Deploy. My deployment folder is created.

Pull the Package Over
3:25-4:15
My next step is to pull that deployment folder over to the server where the role is currently. I can use map drive, I can use external disk -- any way I can get that folder over to the source computer, that's what I need to do.
We're going to go take a look at our Windows Server 2008 R2 server. Now that we've packaged up our migration tools, we need to pull them over to the old server. I'm going to map a drive to my Server 2012 C: drive. I like to do it in the command prompt, where if there are any errors, it will prompt me.
I'm logged in as an Administrator, so no trouble. I've now got a P: drive that links to the C: drive of my 2012 Server. We're going to Copy our Deploy folder and Paste it into the local drive C.

Register the Tools
4:16-4:41
Once we've brought the folder over, we need to register the tools. You open up that folder, open up the subfolder, and you want to double-click SmigDeploy. As soon as it comes back with a PowerShell, then you're all set. You're good to go.
If you need to reopen that window, now it's in Administrative Tools. I can just click on the PowerShell icon here.

Store Features in a Variable

4:42-5:57

The first thing we're going to do is find out what features are installed in this server, so that we can capture them to a file and migrate it over. We're going to store those features in a variable. Variables always start with a dollar sign and they are case-sensitive. We're going to retrieve information that always starts with a Get. We're going to get all the migratable server features and we're storing it in a variable called $c. If you ever want to see the contents of a variable, you just type the name of the variable and it will show you what's in there. This particular server has BranchCache and DHCP installed, and that's what we're going to be capturing.

We're going to export my migration server settings. The features that I want are in that variable, $c, and the path where I want to store it is c:\store. -verbose just says if anything goes wrong, give me the maximum amount of information so that I can troubleshoot. It prompts me for a password. Once it clicks Data, we have successfully packaged up those features and they're sitting in a file in that folder. There's my store folder.

Pull the Package Over

5:58-8:24

Now I've got to get this folder back to the new server. We'll go back to 2012. On the 2012 server, I'm going to do pretty much the same thing to pull that file over.

I map to the C: drive of my DHCP server with Windows 2008. Then I can go in there, there's my store folder and I'm going to Copy that to the C: drive of the server I'm migrating to. We've pretty much got to do the same thing over here. We're going to capture everything in that file to a variable and then process the variable to install the features. I want you to take a look right now. You can see we've got Active Directory and DNS running on this server. Observe, there's no DHCP. When we're done, we are going to be able to get access to DHCP, and we're going to see it show up right on this menu.

I'm going to click on my Migration Server Tools. There's my prompt. I'm going to use the same variable. This time we're retrieving it from the file that's in that c:\store. These are not case-sensitive, by the way. I'm typing in mixed case, just so it makes it easier for you to read the command. Give the password to open the file. It's taken those features, stored it in the variable. Again, if you'd like to see the contents of the variable, just type the name of the variable, and you see it come back.

Whereas on the source server we exported, now we're looking to import. We're taking all the features that are in that variable. If we need any extra information it's in c:\store and -verbose, should something go wrong, I'm going to get my maximum amount of information back. Enter the password for my file. OK, it imported DHCP and installed that feature, but it found out that, in fact, this Windows Server 2012 server is already running BranchCache, so it didn't need to install it.

Verify Installation

8:25-9:21

We can verify that the feature got installed by looking in Server Manager. You can see right there, it just popped up. If, for some reason, the service doesn't start, you can right-click it, Start Services.

That's how we can migrate roles from one server to another. This is pretty generic migration. We still would need to install the DHCP Management Tools. For each specific role that you need to migrate, there's going to be a different set of procedures to go through This was fairly easy. Let's say we were migrating Active Directory Federation Services. This is quite a long series of steps. For each server, you need to do your research. Find out the roles and features installed on that server and see if you can upgrade. See if you can migrate.

How to Upgrade Server Core

9:22-11:23

Before we go, I'm also going to show you how to upgrade Server Core. Server Core is a little bit interesting, because you can't upgrade a Server Core 2008, 2008 R2 to 2012 with a GUI. What you can do is upgrade the Server Core to the 2012 Server Core and then install the GUI. That's the best way to do that.

We have our Windows Server 2008 Server Core. I can check that just by hitting Ctrl-Alt-Delete. You can see it is 2008 R2. I've got my Windows Server 2012 DVD in the DVD-ROM drive, so we'll head over there and run setup. When you do an in-place upgrade, you want to run Setup from inside the old operating system. We'll Install now. I can go online to get updates, we're not going to worry about it.

We need to make sure that we upgrade to a Server Core installation. If I pick Server with a GUI, in a couple of screens when I say upgrade, it's going to check and tell me, no, you can't do that. We'll pick Server Core. I agree to the license terms and then, we're going to pick Upgrade, which is going to keep all my files, settings, and applications. It comes up, says, we've found one thing that might not work properly. I'm not terribly worried about it. I'm just going to hit Next. We're doing an in-place upgrade of Server Core 2008 to Server Core 2012.

Now that we've upgraded our Windows Server 2008 R2 Server Core to a Windows Server 2012 Server Core, the next step would be to go into PowerShell, install the GUI Server Shell, and we'll be up to a full version of Windows Server 2012. That's the best way to upgrade Server Core.

2.3.3 Features on Demand

Features on Demand
0:00-0:47
Let's talk about Features on Demand. This is kind of an interesting new technology coming in with Windows Server 2012. The idea is this; you have different roles or features that can be installed on the server. Let's say that you know for sure you're not going to install DNS. Well, the source files needed to install DNS are still on that server. It's one of the great changes that they made starting in Windows Server 2008 is, they make a copy of all the files that will be needed at any point in time on the hard drive of the server or the client, so that you don't get prompted to insert the installation disk. People that had been around since the beginning of Windows all the way up to Windows Server 2003 are very familiar with the "please insert the i386 folder". No more i386 folder, and no more need for the source disk, if you want to install a feature.

Removing Features
0:48-1:14
Features on Demand allows you to go in and remove those files from the hard drive. The presumption being, if you know for a fact you're not going to use this feature, then you don't need to have the files sitting around on the hard drive taking up space. If you do remove that feature, then you can go through, and it's not like you've given up the option to install it. If you do try to install the feature, Windows will prompt you for where to get those files from, since you've removed them from the hard drive.

Remove Role and Feature Files
1:15-1:37
Features on Demand allow us to remove the role and feature files, with the main objective being to conserve disk space. We can do that on a server that's running. We can even do that on a Windows image--a WIM image--which is Microsoft's image format--or an offline VHD. Any type of place that we have Windows Server 2012, we can remove those files, if we need to, to save space.

Install Roles and Features
1:38-1:53
Once you remove the files, of course, you're going to need to get them back if you do need to install those roles. You can install the roles and features from remote locations. You can have a centralized location that anybody who needs to install a role on a server can pull from, just by popping in the installation media.

When Files Are Not Available
1:54-2:31
When the files are not available, they can be provided by...they have something called the side-by-side feature store, which is really just a shared folder on the network.
The side-by-side feature store, really just a shared folder on the network. I'm going to go through and I have a slide telling you exactly how to do that. You can also get those from the Windows Update, which is kind of cool, so if you did remove the features, using Features on Demand, and you don't have a copy, you can just go up to Windows Update and grab them, or you can put in the installation media. This is going to take bandwidth. That's the only reason why you wouldn't want to do it. Other than that, there's no problem with that. This would be a little bit faster because you don't have to download it, and it won't cause any traffic over the WAN link.

Default Search
2:32-3:10
By default, when those files are not there, it does have an algorithm that it uses to search. It could search a location specified by the user. When I give you the PowerShell command, I'm going to show you how to put that in.
There's also a Group Policy setting that you can use to give it a spot that it can go to look for those files. That way instead of having to set it on each and every server, or specify the location when you run your PowerShell command, you can just have a GPO that hits all your servers and says, "Hey, if any of you guys are missing files, here's a central place where you can go". Or again, the last thing it's going to do is hit the Windows Update. It's going to go through these two first, and then Windows Update is sort of a last choice.

Side-by-side Feature Store
3:11-4:04
The Side-by-side feature store, I think this is great that they have this entire name, the Side-by-side Feature Store. It's just a really fancy term for a share where you stick these files. Of course, if you're going to create one of these, the first thing you do is make a Shared folder. It's important, if you make that shared folder, either to give everybody the ability to read it, or at the very least, the computer accounts needs the ability to read it, so that the computer can hit that share and pull down the files that it needs. If it doesn't have the correct permissions, of course, it's not going to be able to pull it down. Once you've made the folder, you've given the correct permissions, you would go into the DVD on the Windows Server 2012, or even Windows 8 DVD. There will be a folder called sources. In there is

another folder called \sxs. You're going to copy that entire folder into this shared folder. These are all the source files you need for any of the features or roles that might have been removed using Features on Demand.

Remove Features

4:05-4:35

In order to remove the features, it's going to be a PowerShell. Make sure you use an Elevated Command Prompt, which means you're going to Run as Administrator. Here's your main command. Here's your verb, Uninstall. The noun, what you are uninstalling is a WindowsFeature. You would give the -Name of the feature -- unfortunately, you remove them one by one. You could give the -ComputerName if it's a different computer. Then, make sure you have the -Remove, because that's what we're going to do to get rid of it. If it was a WIM file, then you would use a command called DISM.

Install Features

4:36-5:09

If you need to install them, of course, it's going to be Install-WindowsFeature. There's our verb and our noun. You've got to give the Name of the feature. If it's not this computer, we can give the computer name that we're going to be doing it on.

Here's that part I talked about earlier. -ConfigurationFilePath. This would be how you would tell it where to find those. Where I put these x's, you could put in D:\sources\sxs, or it could be where that share is. Wherever it is you're keeping it, this spot here is where you would specify that. Again, if it was a WIM file, you would use DISM.

Summary

5:10-5:31

That's the scoop with Features on Demand. It really just means that you're going to add in the Features on Demand. Of course, you're going to have to remove them first. By default, all the files are on the hard drive, and this would be to conserve space. If you do need to save space, sometimes servers can get very pinched for space. This is something that you can do to get rid of some of those files. Absolutely, you can get those features installed, if you remove them -- very easily, after the fact.

2.3.4 Using Features on Demand

Using Features on Demand
0:00-0:33
In this demo, we're going to take a look at using Features on Demand. We're going to use PowerShell to completely remove the source files that would be used to install this feature. We anticipate we're not going to use this feature, and we want to save some room on the hard drive and just get rid of the source files. In this one I'm going to use DHCP, because it's easy to spell and it's very short. You can see by looking it's not one of the roles that I have installed, so it won't be a problem if we remove it, and then I'll show you how to add it back in.

Removing the Role
0:34-0:43
In PowerShell, to remove the role, we're going to uninstall-windowsfeature. We need to give the -name of the feature that we're going to uninstall; the name here is dhcp.

On a Remote Computer
0:44-0:50
If I were doing this on a remote computer, I would add a switch that would say computer name and give the name of the computer.

On a Local Server
0:51-1:03
Since I'm going to do this on the local server, I don't need to do that, but I'm going to add a -remove which tells PowerShell to completely get rid of the files. Now that I've completely removed the DHCP files from this server, if I did want to install the role, I'd have to provide those files to the server, and there's a number of ways that we can do that.

Installing the Role
1:04-1:20
One way is to specify the location in the command that we'll use to install it.

Side by Side Feature Store
1:21-2:47
Another way to do it is to set up a Side-by-side Feature Store. A Side-by-side Feature Store is simply a shared folder, and you can create this on any server. I'll make a New Folder, it doesn't have to have a particular name. In this particular folder, you have to go into the Windows DVD, inside of \sources, and copy everything from this \sxs folder into the folder that you created. Highlight everything, and I would just simply copy it into FeatureStore, and then I need to make sure that I share this folder. I can share it with specific people; if I want to do everyone read, that's fine. At a minimum, the computer accounts need to have access to that. If I am in a domain, I could do domain computers, whatever group you want to add, but they need at least Read to this folder once it's shared out. I can specify the Side-by-side Feature Store in the command itself.

Group Policy
2:48-3:41
I can also go through and set up a Group Policy to affect all of my servers to tell them all to use the same spot. This is inside of \Policies on the \Computer Configuration side. I'm going to open up \Administrative Templates and click on \System. You don't have to. I click over on Standard because I can see it a little bit better, and if you scroll down midway through here you will see "Specify settings for optional component installation and component repair". That's the Group Policy that you need to turn on. I'll double click it, and if I enable it, I would then give the UNC path of my feature store that I've created. I also can specify whether or not my computers are going to be allowed to talk to Windows Update, because the first place it's going to search is wherever the user specifies, second place, it would use this GPO -- where I would specify my feature store, and then its next spot would be to go and look at Windows Update.

Windows Update
3:42-4:57
If I don't want it to look at Windows Update I can say, "Never attempt to download those files from Windows Update". Or if I have a Windows Server Update Services in my environment, but I don't want the servers to contact WSUS, I'd rather have them go directly to Microsoft, I can enable that as well. I can set up where they'll get the files from using Group Policy. In our case, we're simply going to specify where they should get the file from, so I'm going to install a Windows feature. The name of the feature is DHCP, the source is going to be my DVD drive that we just looked at, and we should be good to go. Now DHCP has been installed. Notice, too, that a restart is not needed, but if it was, you could build that right into your command just by adding a restart switch.

Installing Management Tools

4:58-5:25

The other thing you also might want to note is that if you use this on a Server Core machine, it's not going to install any management tools by default; so there's also a switch that you can add that says -includemanagementtools, which would cause it to install the management tools as well. But most Server Core machines you might be administering via remote, so it might not be an issue. That's just a little demonstration on how we can use features on demand.

2.3.5 Server Roles Facts

Functionality and services are added to a server using the Add Roles and Features Wizard. The software is categorized as follows:

- A *role* is a set of software features that provides a specific server function. Examples of roles include DNS Server, DHCP Server, File and Storage Services, and Print and Document Services.
- *Role services* are specific programs that provide the functions of a role. Some roles, like DNS Server, have a single role service. Other roles, like Print and Document Services, have multiple role services such as the Distributed Scan Server and Internet Printing. Role services are sub-components of a role.
- A *feature* is a software program that is not directly related to a server role, but which adds functionality to the entire server. Features include management tools, communication protocols or clients, and clustering support.

All roles, role services and features are added using the wizard. You remove a role using the Add Roles and Features Wizard.

Roles in Windows Server 2012 include:

Role	Description
Active Directory Domain Services (AD DS)	AD DS is a distributed database that stores and manages information about network resources such as users, computers, and printers. The AD DS role: • Helps administrators securely manage information. • Facilitates resource sharing and collaboration between users. • Is required for directory-enabled applications such as Microsoft Exchange Server and for applying Windows Server technologies, such as Group Policy.
Active Directory Certificate Services (AD CS)	AD CS is an identity and access control feature that creates and manages public key certificates used in software security systems. The AD CS role: • Provides customizable services for creating and managing public key certificates. • Enhances security by binding the identity of a person, device, or service to a corresponding private key. • Includes features that allow you to manage certificate enrollment and revocation in a variety of scalable environments.
DNS Server	The DNS service maps IP addresses to logical hostnames. DNS servers provide name resolution services, providing IP addresses for known hostnames or hostnames for known IP addresses. Beginning with Windows Server 2008, the DNS service provides support for IPv6 addresses.

DHCP Server	The DHCP service provides IP addresses and other IP configuration information for network hosts. Host computers contact the DHCP server at startup to obtain IP address, default gateway, DNS server, and other configuration information. Beginning with Windows Server 2008, the DHCP service supports IPv6 addressing and configuration information.
File and Storage Services	File and Storage Services includes technologies that help you set up and manage one or more file servers. This role is useful when users need access to the same files and applications, or if centralized backup and file management are important to the organization. New features/functionality includes: • Storage Spaces and storage pools enable you to virtualize storage. • Unified remote management of File and Storage Services in Server Manager enables you to remotely manage multiple file servers. • Windows PowerShell cmdlets for File and Storage Services allow you to perform the majority of administration tasks for file and storage servers. The File and Storage Services role is installed by default on Windows Server 2012.
Hyper-V	The Windows hypervisor provides the layer of software necessary for the installation of virtual guest operating systems.
Print and Document Services	The Print and Document Services role provides the print management console that allows you to manage printers on multiple servers. Beginning with Windows Server 2008, you can also publish printers in Active Directory, thereby creating printing objects on client computers automatically for shared or network printers.
Network Policy and Access Services	Network Policy and Access Services, formerly Network Access Protection (NAP), are a collection of components that allow administrators to regulate network access or communication based on a computer's compliance with health requirement policies. Network Policy and Access Services give you the ability to restrict access for non-compliant computers as well as to provide access to updates or health update resources to allow computers to become compliant.
Web Server (IIS)	Web Server (IIS) is the Web server service. Use IIS to host internal and external Web sites or services that communicate using HTTP and to provide support for ASP.NET applications accessed through a Web browser. IIS is also used by many other roles to provide Web-based administration or access.
Windows Deployment Services (WDS)	Windows Deployment Services (WDS) is a disk imaging solution that you can use for remote deployment and automated installation of Windows

	Server 2012, Windows 8, and earlier versions of Microsoft operating systems.
Windows Server Update Server	The WSUS server allows administrators to manage and distribute updates through a management console. A WSUS server can also be used to update other WSUS servers within the organization.

Role services in Windows Server 2012 include:

Services	Description
Distributed Scan Server	The Distributed Scan Server provides services that route documents scanned on network scanners. The Distributed Scan Server includes the Scan Management snap-in for configuring and managing network scanners.
Network Policy Server (NPS)	Network Policy Server (NPS) allows you to centrally manage network access through a variety of network access servers such as VPN servers, 802.1X Ethernet switches, and RADIUS-compliant 802.1X wireless access points. • NPS contains Network Access Protection components. • NPS allows you to use Protected Extensible Authentication Protocol (PEAP)-MS-CHAP2 for secure password authentication on wireless connections.
Host Credential Authorization Protocol (HCAP)	Host Credential Authorization Protocol (HCAP) allows you to perform client health evaluations and authorization of Cisco 802.1X access clients on networks integrating NPS and NAP with Cisco Network Access Control Server.
Health Registration Authority (HRA)	Health Registration Authority (HRA) issues health certificates to clients on networks using NAP IPsec enforcement.

Features in Windows Server 2012 include:

Features	Description
Failover Clustering	Failover clusters provide high availability and scalability to servers including server applications such as Microsoft Exchange Server, Hyper-V, and Microsoft SQL Server. The server applications can run on physical servers or virtual

	machines. Failover clusters can scale to 64 physical nodes and to 8,000 virtual machines.
Group Policy	Group Policy allows you to specify configurations for users and computers through Group Policy settings. • Whenever the Group Policy Management Console (GPMC) is installed, the Windows PowerShell module is also installed. • If you install the Remote Server Administration Tools pack, the latest Windows PowerShell cmdlets for Group Policy are also installed.
Network Load Balancing	By managing two or more servers as a single virtual cluster, Network Load Balancing (NLB) enhances the availability and scalability of Internet server applications such as those used on web, FTP, firewall, proxy, virtual private network (VPN), and other mission-critical servers. • NLB allows all of the computers in the cluster to be addressed by the same set of IP addresses. • NLB maintains a set of unique, dedicated IP addresses for each host.
BitLocker Drive Encryption	BitLocker Drive Encryption is a security feature that protects a server by encrypting the operating system volume and verifying the integrity of other startup components. BitLocker is also called full volume encryption.
Remote Assistance	Remote Assistance enables a support person to offer assistance or reply to requests for assistance from desktop users. With Remote Assistance, the helper can connect to the computer desktop to watch or perform tasks to troubleshoot and correct desktop problems.
SMTP Server	The Simple Mail Transfer Protocol (SMTP) is used for transferring mail between e-mail systems and some e-mail clients. Add the SMTP Server feature to add e-mail support to other server roles such as IIS.
Windows Server Backup	Windows Server Backup provides backup and recovery for Windows Server 2008 and Windows Server 2012. It replaces the **NTbackup.exe** backup utility in previous Windows versions. Windows Server Backup allows you to manage backup and recovery from either the command line or the Windows Server Backup console snap-in.

When Windows Server 2012 is installed, the source files for all server programs, including all roles and features, are installed on the server. Features on Demand is a new feature in Windows Server 2012 that allows you to remove the source files of unneeded roles and features in order to conserve disk space. Features on Demand also allows you to re-install source files for roles and features that may have been removed. With Features on Demand, you can:

- Add or remove role and feature source files on a remote computer.
- Add feature files to or remove feature files from Windows image (WIM) files or offline virtual hard disks (VHDs) to create a custom Windows Server 2012 configuration.
- Remove feature files from running physical or virtual computers.
- Obtain the files to install from:
    - A shared folder that contains feature files and is available to the computer
    - Windows Update
    - Installation media

2.4 Server Core

As you study this section, answer the following questions:

- How does a Minimal Server Interface installation differ from a Server Core installation?
- What are the benefits of a Server Core installation?
- Which utility do you use to make changes to Server Core settings?
- Which PowerShell command converts a Server with a GUI installation to a Server Core installation?
- Which four tasks should you perform immediately after a Server Core installation?
- Which command joins the server to the domain after the Server Core installation?

After finishing this section, you should be able to complete the following tasks:

- Install and configure a Windows Server 2012 Server Core system.
- Convert a Windows Server 2012 Server Core installation to a full installation.
- Convert a Windows Server 2012 full installation to a Server Core installation.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.

This section covers the following 70-410 exam objectives:

- 101 Install servers.
  This objective may include but is not limited to:
    Plan for a server installation
    Install Server Core
- 102 Configure servers.
  This objective may include but is not limited to:
    Configure Server Core
    Convert Server Core to/from full GUI
- 203 Configure servers for remote management.
  This objective may include but is not limited to:
    Configure Server Core

2.4.1 Server Core

Server Core
0:00-0:41
Now we're going to talk about Server Core. Server Core is an installation option. When you go through the wizard to install Windows Server 2012, you're going to see that you can install the full server with the graphical user interface, or you can install Server Core. Server Core, then, doesn't have a Graphical User Interface. When I first saw Server Core, back in Windows Server 2008, it was kind of a culture shock. I could see this screen, and floating around in the middle was a command prompt, and that's it. No taskbar, no nothing, and it's been quite a number of years since I've only had a command prompt to do. Of course, silly me, the first thing I did was close the command prompt, but we'll show you in the demo how to get that back, if you need to.

Benefits of Server Core
0:42-1:19
There are plenty of benefits to Server Core that you can get, so you might be thinking, "Why would I want something that's just a command line?" It sounds like it will be harder to support-- not at all. You can use the Graphical User Interface on a full server and redirect the snap-in. You can also install the Remote Server Administrative Tools, RSAT, which is a free download from Microsoft. With Server Core, you're not committing to only using the command prompt. It's just that there's no GUI running on that server. There are benefits to doing that. Let's take a look at some of those benefits.

Reduced System Requirements
1:20-2:05
The first thing that we would say in favor of Server Core is that there are reduced system requirements. Specifically, it doesn't require as much RAM, because it's not actually going to produce that graphical user interface, and the operating system files use less disk space. It's ironic that there is a push with Server Core and Features on Demand to save disk space, while at the same time, Microsoft has a philosophy that disk space has become really cheap, so if you're doing backups, you can do backups on a removable drive. Sometimes, there's a little bit of a conflict there. Always plan for adequate disk space. Don't be thinking, "Well, I'll install Server Core, and then I can just buy a really small hard drive." Hard drive space and RAM, I would say it's like money. I can never have too much of any of those things.

Fewer Components to Troubleshoot
2:06-2:20
Since we're not installing the Graphical User Interface, there would be fewer components to troubleshoot, so it's not going to be a video problem or a problem with the snap-in, because the snap-ins simply aren't there. Command Line Interface is there, but the GUI snap-ins are not there.

Reduced Servicing Requirements
2:21-2:36
We also should have reduced servicing requirements. We wouldn't have as many patches and updates to install, because we don't have as much running. A lot of times, it's the Graphical User Interface that creates a security issue, or creates an opening for hackers.

Reduced Attack Surface
2:37-3:03
That leads us to our next bullet, which is that it has a reduced attack surface. The big thing is, it doesn't have Internet Explorer. It doesn't have Windows Explorer, so we have less open ports, less services, so hackers are not going to have as much to target, because we just have less things being installed. It's something like more than 30 services are not installed with Server Core, than are installed with the full server installation.

Summary
3:04-3:18
Server Core is an installation option that we can use. It can make the server more stable, less vulnerable to attacks, requires less disk space, and it's going to run a little bit better. We don't need to interact with the GUI, then we can take it off.

2.4.2 Configuring Server Core

Configuring Server Core
0:00-0:28
In this video, we're going to talk about configuring Server Core. For the most part, we want to configure Server Core by remotely managing it from a server that has a GUI. When you first install a server, there are certain things that need to be done before that can happen. Specifically, we need to make sure it has the correct name, make sure that time and time zone is set correctly, give it a static IP address, if it needs one, and join the domain. We're going to look at how to do that from the command prompt in Server Core.

Correct Name
0:29-1:28
When I install Windows Server 2012, it's going to create a random server name. You can see what the server name is, or the computer name of any computer, from a command prompt using the hostname command.
We can see that it's been given some big, long name. Probably not what we want to maintain as the name for this server, so we're going to change it. I'm going to need to use the existing name in the command to change the name. I don't know about you, but I tend to make lots of typos. What I'm going to do is copy this to the clipboard, so I can just paste it in. You can copy anything from a command prompt by right-clicking the command prompt, and choosing Mark, and that lets you highlight. I'm going to highlight the existing name. Once it's highlighted, hit Enter on the keyboard, and that will put that information into the clipboard, so I can paste it in later. As soon as I hit Enter, the highlighting goes away, and it's in the clipboard.

Command to Change Computer Name
1:29-1:52
Now, we'll type up the command that will change the computer name. My command is, netdom renamecomputer. There is the existing name, and the newname will be Member1. It comes up and it warns me, changing the name may impact services. That's exactly why we want to get the correct name before we start installing anything on the server.

Reboot the Computer
1:53-2:24
Now, it tells me that the computer is going to need to be rebooted to complete the operation. You can reboot the computer using the shutdown /r command. I'm adding a /t 0, which tells it that the amount of time to wait before it reboots is no time at all. If you just use /r, that would work fine, too. It would delay about a minute, and then it would reboot.
Now that my server has rebooted, I can see that the name has been successfully changed.

Set up Time Zone
2:25-3:00
Now that I have named my computer, the next thing I want to do is make sure that my time zone is set up. Computers are very sensitive about time, because they don't want anybody to be able to record packets, and then play them back. It's called a replay attack. The times have to be synchronized throughout the domain. I'm just going to go ahead and check the time and the time zone on this server. You can see this opens up the standard Date and Time interface, and I can make any adjustments that I need to.
If you need to change the time from the command prompt, you can actually just use the time command. Now, I have the correct computer name. I have the correct date and time.

Set a Static IP Address
3:01-4:59
What I'm going to do is go ahead and set a static IP address. The first thing I need to do is find out the index number of the network interface. Anything with TCP/IP from the command line is done using the netsh command. There are two ways to do this command. One is interactively, and the other is to just type out the entire command at the command prompt.
I'm going to show you a brief demonstration of both, and then we'll just stick with typing the entire command. You can see my prompt changed to netsh. I'm now in interactive, and I can build by changing my context, so that when I actually type the command, it will be a little bit shorter, less chance for typos. Now I'm looking at the interfaces and I can see the index number of the interface I'm interested in is 12. I can do the same command as just one long command at the command prompt. You can see, I get the same output. Now that I know I'm working with adapter number 12, I'm going to type in the command that will set a static IP address for this adaptor.
If all I wanted was DHCP, the computer is set to use that automatically, so I'd be all set. I could just continue. I'm going to go ahead and show you how to set a static IP address, because many servers do require a static IP address. I have netsh interface, ipv4. I'm going to set an address, name= and I use the index number of the adaptor that I want to change. The source=static, which means it's a static address. It's going to get IP address 192.168.1.51,

subnet mask of 255.255.255.0, and a default gateway of 192.168.1.1. I can use ipconfig to make sure that that went off successfully. There's my address.

Joining the Domain

5:00-6:12

In order to be able to join the domain, I need to make sure that this computer is going to use a DNS server that has the records for the domain. I can see right now, I'm not really sure what those addresses are, but they're probably not the ones I want. I need to add a DNS server that's going to allow me to contact the domain. My command is netsh interface ipv4 add dnsserver name=12 specifies the Ethernet adapter I'm going to modify. The address of my DNS server is 192.168.1.40. index=1 means it's going to be the primary DNS server. If I already had another primary DNS server configured, setting the index of this one-to-one would just bump the other one down to being secondary. Even if I want the original primary to become secondary, I still only need the one command. I can make sure that my command was successful by doing an ipconfig /all. There's my DNS server.

Pinging the Domain

6:13-6:29

You also can make sure everything is all set to go to join the domain by pinging the domain. The fact that I get a reply from pinging the domain, tells me DNS is working okay. The domain controller is up, I can find the domain controller.

Join Server Core Serer to the Domain

6:30-7:41

At this point, I'm ready to join my Server Core server to the domain. I'm going to be joining Member1 to the domain, the user account, and it is userd, not user. The user account I'll be using is administrator; passwordd. Notice there's an extra d.

I put a star so it will prompt me to type in the password. I could just type the password there, but it would show in plain text. I'm going to use a star, so that if somebody were watching me, they wouldn't know what the administrator password is. One thing that bothers a lot of people is I don't see any asterisks, so they think nothing is happening. Just have faith. Type the correct password, and hit Enter. I've joined the domain. As soon as I restart, I'll be a member. Now that I've rebooted, I'm a member of the domain. I wouldn't recommend memorizing netsh, or netdom commands, but I just wanted you to see the way to interact with the computer through the command prompt.

Shortcut

7:42-8:15

There is a little bit of a shortcut. I would never show you that first, because then you wouldn't be quite as interested in my other commands. If I type sconfig, this is the closest you're going to get to a GUI on a Server Core machine. It shows me that I am a member of the domain. The name of the computer is Member1. I can go through. I can set my Network Settings, Date and Time. I can actually use this interface to make any changes that I need. It's good to have an idea of both ways to do it. That's how we configure Server Core from the command line.

2.4.3 Server Core Conversion

Server Core Conversion
0:00-0:24
We're going to talk about Windows Server Core conversion. Server Core was an installation option that came in with Windows Server 2008. In 2008 R2 and 2008, you either had Server Core or you had a server with a GUI. Once you made that decision, you were kind of stuck. New with Server 2012, you can switch between Server Core and between this full GUI, and you can actually stop halfway in the middle.

Three Modes.
0:25-0:25
There are three modes.

Server Core
0:26-0:45
We'll start with the most secure. Server Core, there's no GUI shell. There's no GUI management tools. The shell would be the taskbar, the Start menu. The management tools would be like, Server Manager, any of the snap-ins. This would be the most secure and the most stable. It's the most secure and stable because it's not going to need as much patching -- it's not running as many services.

Minimal Server Interface
0:46-1:16
In between that, there's actually something they call the Minimal Server Interface. This has most of the GUI management tools, but not the shell, so no taskbar, no Start menu, and not Internet Explorer, which is great. You definitely don't want LAN administrators sitting in the server room surfing the internet. It might be entertaining for them, but great way to pick up viruses and spyware. It's going to be more secure and more stable than the full GUI, but it's kind of half and half. It's really weird to look at it. You'll see that in the demo because you have Server Core, but then, Server Manager floating in space above it.

Full GUI
1:17-1:36
Then, at the far end of the extreme, we have the Full GUI, which is going to be the easiest to use for administration, because you've got all your tools. You've got your Start Menu that looks like normal Windows Server 2012, but it's going to be the most difficult. It doesn't mean it's difficult, but relatively speaking, the most difficult to support, because it's going to need the most patches, and it's going to have the most services running and ports open.

Server Core Conversion
1:37-1:47
Server Core Conversion can be done either using PowerShell or Server Manager. We're going to do some basic PowerShell information, and then we'll look at the specific things you install and uninstall.

PowerShell
1:48-3:08
Before switching using PowerShell-- this is going to be important if you're in Server Core and you're trying to get over to the GUI--you have to go in and do Import-Module ServerManager. That will bring in the tools that you need so that you can install and uninstall the features. To get into PowerShell on Server Core, because all you have is a command prompt, you just type PowerShell and hit Enter, and that puts you into PowerShell. It's actually pretty simple. Then, to install a feature, you use the Cmdlet (anything in PowerShell is a Cmdlet) install-windowsfeature. Again, PowerShell's not case sensitive. To uninstall it, very simple, uninstall-windowsfeature. For the Server Core conversion using PowerShell, if you want the full server installation, you can install the server GUI shell. That's going to pull in the server GUI management infrastructure. When you see it in Control Panel, you'll see two options, the shell and the infrastructure. If you want the minimal server, you have to go to the full server and then uninstall the shell. That will leave behind the management infrastructure or the tools without the shell, that taskbar and Internet Explorer. If you want the Server Core installation, the easiest way to do it is to just uninstall the management infrastructure, and when you take that out, it's going to pull the shell out with it. It's kind of neat. For each one of them you can run one command and it'll get you where you need to go.

Remote Role Deployment
3:09-3:46
In PowerShell, the management tools and snap-ins for a role are not included by default, so if we install a role in Server Core, we're not going to get the management tools. If we want them, you can add the -IncludeManagementTools switch to the PowerShell cmdlet, and what that will do, essentially, is take your Server Core and switch it over to a minimal shell option. That will have the management tools running. You'll be prompted to do that and it will become minimal. It's kind of an interesting mix, and you've got to really be aware of which

deployment type you'll have at any given time. You want to be sure you're doing it purposefully. It doesn't matter what you choose, as long as you've actually chosen it.

Server Core Conversion - Server Manager

3:47-4:26

In Server Manager, for the full server installation, we need both of them. It's going to be inside the User Interfaces and Infrastructure Category. You'd want to check both Graphical Management Tools and Infrastructure and Server Graphical Shell, because the full server has both of those. For minimal, we would just uninstall the graphical shell. Leave the management tools and infrastructure checked, and we'll have the minimal server installation. Finally, for Server Core installation, uninstall both objects. Go into this category and basically uncheck everything. That's how you can convert between Server Core and the full GUI installation. It means that you don't have to make up your mind until you are ready to make up your mind.

2.4.4 Converting Server Core

Converting Server Core
0:00-0:09
In this demo, we're going to take a look at converting between the full Graphical User Interface installation of Server 2012 and the Server Core installation.

Full GUI and Server Core Comparison
0:10-0:28
When you install Server 2012, we have a choice. We can go with the full GUI, which is what we're looking at right now. Here we have the taskbar down at the bottom of the screen. We've got all of our snap-ins; Server Manager available; we have Internet Explorer if we need it.
In Server Core, we wouldn't have any of those things.

Minimal Server
0:29-1:16
Now there's actually a third option, that sort of splits the difference, and they tend to call that the minimal server. With a minimal server, I have my snap-in utilities, such as Server Manager or my management consoles, but what I don't have is the shell. I wouldn't have the taskbar; I wouldn't have Internet Explorer. That's the minimal option.
We can switch between these using either Server Manager, or we can use PowerShell, and I'm going to show you both. To get to the minimal server, you've got to start out with the full GUI. So we're going to start here with the full GUI; we'll roll back to minimal, and then we'll roll completely back to Server Core, and then we'll come back around to the full GUI.
In my Add Roles and Features Wizard, I want to go ahead and get into Features.

Converting from Full GUI to Minimal Server with Server Manager
1:17-1:29
I'm going to manage this particular server; I'm not worried about roles; the particular features that we're looking at are down here, User Interfaces, and Infrastructure and you can see that it's installed.

User Interfaces and Infrastructure
1:30-2:41
There's two parts to this. One is the Graphical Management Tools and Infrastructure; the other is the Server Graphical Shell.
That shell is the taskbar and Internet Explorer. were to uncheck this, then I would roll back to my minimal server. If I uncheck both of them, that will roll me back to Server Core. You can see I can't uncheck these. That's because I'm in the Add Roles and Features Wizard. If I really wanted to uninstall them, when I start this wizard, you need to make sure that you Start the Remove Roles and Features Wizard, which is very easy to overlook.
The same exact wizard, but now I'm removing. Now I'd be able to uncheck them. I'm not going to do that, because I want to show you how to do it with PowerShell. Now, because I want to go back to minimal server, I'm actually only going to uninstall the shell.

Converting from Full GUI to Minimal Server with PowerShell
2:42-3:20
My command is uninstall-windowsfeature (always a verb-noun space) server-gui-shell.
Now that it's finished, I have to restart the server to finish the removal process; could have added a -restart, or I'm actually just going to restart it from the command line using the shutdown command. Now that my server is rebooted, I should have a minimal server installation.

Minimal Server Installation
3:21-3:46
You can see the minimal server installation has Server Manager; it's got all the tools that I'm used to having. My Tools menu -- but you can clearly see there's no taskbar down at the bottom of the screen; there's no Internet Explorer. The only interface I have is the command prompt -- very similar to Server Core.

Why You Would Want Less GUI
3:47-4:28
Why would you want this? The more of the GUI that you have installed, the more patches you tend to need; the more services you have running, which could be attacked either by hackers, or by viruses, or spyware. Many a great server has been brought down by a LAN administrator surfing the internet, or even if they're just looking for technical information and they happen to install a virus or spyware using Internet Explorer.
This kind of gives us the functionality of the GUI; a little bit less risky, but it's also going to be a little bit more difficult to use, because I don't have that taskbar, Start Menu, Internet Explorer, that type of thing.

Converting from Minimal Server to Server Core

4:29-4:54

The next step down from here would be Server Core, where we won't have anything of the GUI. In order to do that, we're going to go into PowerShell; I can get into PowerShell from the command prompt just by typing powershell. I could also have gotten PowerShell from the Tools menu, but this is just as good. Now I'm going to uninstall the management tools.

Converting from Full GUI to Server Core

4:55-5:30

If you uninstall the management tools from the full GUI, it will take the shell with it. I could have skipped this minimal server and gone straight to Server Core, just by uninstalling management tools. We just did it in two steps so that you could see the minimal server.

Now that it's finished we just reboot the server and we'll be at Server Core. Now that it's rebooted, we should be at Server Core. Let's go take a look. That's all we're going to see in Server Core, is just the command prompts.

Opening Command Prompt in Server Core

5:31-5:58

One of the things I accidentally did when I first saw Server Core was to close the command prompt, and then I said, how am I going to interact with the server.

Quick tip: you hit Ctrl Alt Delete, go to Task Manager, go up under File, Run new task, and the new task I want to run is command prompt, and that's how you get your command prompt back.

Converting from Server Core to Full GUI or Minimal Server

5:59-7:02

Now that we've got Server Core, we're going to roll back to the full GUI. I don't have to do it in two steps the way we did it with the removal if I just install the shell that's going to pull the management tools in with it.

I can go right back to the full GUI in one step, or if I so desire, I could install that server\gui\management\infa; that'd bring me back to minimal server, then install the shell; sounds like a lot of work to me, let's just go install the shell and go right back to the full GUI. I've got to go into PowerShell in order to do this. I'll type my command. And this time, we'll be clever and add a restart so that it will automatically reboot.

Now that we've rebooted, we should be back to our full GUI. Let's take a look. There's our shell come up; taskbar, and there's Server Manager.

Summary

7:03-7:43

That's how you switch between the full GUI, the minimal installation, and Server Core.

The advantages of Server Core: it's more secure, it's more stable, but generally, we end up administering these machines remotely from a server that has the full GUI. Even though you can choose to change between the installation options, it's a good idea to decide what kind of a server this particular server is going to be and not be installing/ uninstalling the GUI frequently. It should be something that you do seldom, if ever.

It is nice to have the choice if you need to add it in or take it out.

2.4.5 Server Core Facts

When you install Windows Server 2012, you can choose between Server Core Installation, Server with a GUI, and a third, intermediate option, Minimal Server Interface. The Server with a GUI option is the equivalent of the Full installation option available in Windows Server 2008 R2. Server core is a minimal server installation option which provides a low-maintenance environment capable of providing core server roles for computers running on the Windows Serve 2012 operating system.

In Windows Server 2008 R2, the installation option choice between Server Core and Server with a GUI was permanent. In Windows Server 2012, you can switch between a Server Core, Server with a GUI, and Minimal Server Interface. The following table describes the installation options.

Option	Description
Server Core	When you choose the Server Core installation option: • There is no GUI shell. • There are no GUI management tools. • Server Core provides the most secure and stable installation of Windows Server 2012.
Minimal Server Interface	When you choose the Minimal Server Interface option: • There are GUI management tools. • There is not GUI shell. • Internet Explorer is not available. • The Minimal Server Interface is more secure and stable than a Server with a GUI.
Server with a GUI	When you choose the Server with a GUI option: • The GUI management tools provide the easiest methods of administration. • The need for more security and patches increases the difficulty and burden of supporting Windows Server 2012.

Consider the following when deciding which installation option to choose.

- The benefits of Server Core are:
  - Stable environment
  - Reduced system requirements
    - Requires less RAM
    - OS files use less disk space
  - Fewer components to troubleshoot

Reduced servicing requirements
  - Reduced patching
  - Reduced updating

Reduced attack surface
  - Fewer open ports to target
  - Fewer services

- Server Core's minimal structure creates certain limitations, such as:
  Server Core has very little GUI functionality.
  The interface is a command prompt with PowerShell support.
  There is only limited MSI support when used in unattend mode only.
- Server Core supports the following Windows Server 2012 roles:
  Active Directory Certificate Services
  Active Directory Domain Services
  DHCP Server
  DNS Server
  File and Storage Services
  Active Directory Lightweight Directory Services (AD LDS)
  Hyper-V
  Print and Document Services
  Streaming Media Services
  Web Server
  Windows Server Update Server
  Active Directory Rights Management Server
  Routing and Remote Access Server

Server Core also supports many Windows Server 2012 server features.

Perform the following tasks immediately after the Server Core installation:

Task	Process
Name the computer	The process for naming the computer is: 1. Use the **hostname** command to display the generated server name. 2. Use the **netdom renamecomputer** command to rename the computer. 3. Reboot the computer using the **shutdown /r** command. 4. Verify the name has been changed by entering **hostname** and verifying the name change.
Set the time and time zone	As a security measure against replay attacks, the time must be synchronized on servers in the domain. 1. Use the **control timedat.cpl** to display the Date and Time dialog. 2. Adjust the date and time, and time zone.

Assign a static IP address to the server	The process for assigning a static IP address it to use **netsh** commands to: 1. Determine the index number of the adapter. 2. Set the IP address, the subnet mask and the default gateway. 3. Add a DNS server that has addresses for the domain.
Join the server to the domain	The process for joining a server to the domain is: 1. Use a **netdom** command to join the server to the domain. Use an * in the password field to prevent the password from being displayed in plain text on the monitor. 2. Reboot the computer to become a member of the domain.

Enter **sconfig** at a command prompt to display the Server Configuration utility. You can use this interface to make changes to the server settings.

PowerShell commands for switching between server installation options are:

- **Import -Module Server Manager** imports the Server Manager module
- **Powershell** starts PowerShell on Server Core
- **Install-WindowsFeature Server-Gui-Mgmt-Infra** converts Server Core or to Server with a GUI using PowerShell
- **Uninstall-WindowsFeature Server-Gui-Shell** converts Server with a GUI to a Minimal Server Interface
- **Uninstall-WindowsFeature Server-Gui-Mgmt_Infra** converts Server with a GUI to Server Core
- **-IncludeManagementTools** includes management tools when converting Server Core to a Minimal Server Interface. Include this cmdlet with **Install-WindowsFeature Server-Gui-Mgmt-Infra**

Convert from a Server with a GUI installation using the Remove Roles and Features Wizard in Server Manager. The Server with a GUI installation includes the Graphical Management Tools and Infrastructure and the Server Graphical Shell subcomponents of the User Interfaces and Infrastructure feature:

- To convert from Server with a GUI to a Minimal Server installation:
  Leave the Graphical Management Tools and Infrastructure installed
  Uninstall the Server Graphical Shell
- To convert from Server with a GUI to a Server Core installation:
  Uninstall the Graphical Management Tools and Infrastructure
  Uninstall the Server Graphical Shell

2.5 Remote and Offline Servers

As you study this section, answer the following questions:

- How does the remote management of servers increase efficiency?
- What is a server pool?
- How are servers added to a server pool?
- Which roles do not support **Manage As** to change credentials?
- What are the steps required to install and manage roles on a remote server that is not joined to the domain?
- What are the requirements for deploying a role to a VHD file?
- What software is required to use an answer file with Deployment Image Servicing Management (**DISM**)?

After finishing this section, you should be able to complete the following tasks:

- Configure a server for remote management.
- Deploy roles and features to a remote server.
- Deploy roles and features to an offline image.
- Use the WINRM command to manage a remote server

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.

This section covers the following 70-410 exam objectives:

- 102 Configure servers.
  This objective may include but is not limited to:
  Add and remove features in offline images.
  Deploy roles on remote servers.
- 203 Configure servers for remote management.
  This objective may include but is not limited to:
  Configure WinRM
  Configure down-level server management
  Configure servers for day-to-day management tasks
  Configure multi-server management
  Manage non-domain joined servers

2.5.1 Remote Role Deployment

Remote Role Deployment
0:00-0:28
We're going to talk about remote role deployment. This is a really neat feature, and the idea behind it is that in order to deploy a role on a server, I don't actually have to connect up to that server by redirecting the snap-in or by going in and remote desktopping into it. I can actually go in and customize Server Manager so that I can see all my servers in one spot. If I go to install a role or a feature, I can choose exactly which server to install it on.

Server Pool
0:29-0:45
Remote Role Deployment, we actually are going to go ahead and add the remote servers to Server Manager. This is called the Server Pool. If you see the term Server Pool, it just means a bunch of servers that have been added to Server Manager. It's possible that you might need different credentials.

Manage As
0:46-1:23
What you do is, you add the server, and then you click Manage As. You right click it, hit Manage As, and then you can put in different credentials. Between the time you add it and you do the Manage As, you might get an error message, simply because it doesn't have the right credentials. You can ignore that error message. Some roles do not support this and that would Remote Desktop Services, RDS and IP Address Management, IPAM. In that case, you're not going to be able to use Manage As; you would have to actually log on to the server that's running Server Manager with the correct account that has rights on those servers. That's a little bit annoying; probably not the best to do the server pool with those two.

Adding Servers to the Server Pool
1:24-2:00
Adding servers to the Server Pool; really neat, you've got a lot of choices. There's going to be an Active Directory tab, and you can just add servers from the domain, assuming they are all in the domain. There's a DNS tab that will let you add servers by name, or even by IP address. Even though it says DNS, we can still use the IP address to add the server, or there's going to be an Import tab, and we can make a text file. To design the text file, you just have one name or IP address per line, you import the text file, and all those servers will show up. You have to configure several servers identically, you could use the same text file on a couple of different management systems.

Add Workgroup Servers to the Pool
2:01-4:17
Here's where it gets complicated. You have Server Manager in a domain, but you're trying to add a Workgroup Server to the pool. This is a server that's not in the domain. Why would you want to do that? Sometimes in companies, we have public machines. They tend to call them kiosks, so maybe in the cafeteria you have a machine that anybody can use to surf the internet, or maybe there's one in the lobby for guests. For security purposes, you don't want this machine on the domain. All it's going to do is be used to surf the internet. We're not really sure that people are even going to have credentials. Those machines might be in a Workgroup, and yet you may still want to use this technology to manage them from a single interface.

Because these servers are not part of the domain, we have to tell the computer running Server Manager that they are trusted hosts. You simply run this command and where it says <SERVERNAME> here you would put the name of the server that you want to go ahead and add. Make sure you put in the -Concatenate switch and the -Force switch. Both of those are important, and that tells the computer running Server Manager, 'Hey, this is an okay server to talk to.' We also have to tell the server that's being managed, it's okay to talk to the computer running Server Manager. If these two servers are in the same subnet or if the Workgroup server's network connection is set to Private, you don't have to do anything, no change is needed. If that's not the case, you've got to go into the Workgroup server, and in the inbound firewall rules, there's going to be one called Windows Remote Management, HTTP-In, and you would go in and explicitly allow connections from the computer running Server Manager. Once we've done that, if we need to, we can do a Manage As. If you're going to use the local administrator password of the Workgroup computer, that would work great. The problem is, you might not want to give the person using Manage As those those credentials. You could use a different user account that does have administrator privileges, but isn't the one named administrator.

If that's the case, you might have problems with the UAC. If you're going to go with that option, you have to go and run a command to override the UAC by creating a new registry key. And you would actually just run this command on that workgroup computer; it would create a registry key that would override the UAC for that remote connection.

Server Manager on a Workgroup Server

4:18-5:01

It might be that you have that problem in reverse. Maybe you're running Server Manager on the workgroup computer, but for whatever reason you're trying to add computers from the domain to the Server Pool. It's going to be pretty much the same drill.

On the computer that's running Server Manager, we still have to say that the server being managed is a TrustedHosts-- that here's the exact same command we just saw, again, where it's a <SERVERNAME> with the actual name of the server that you're going to be managing, don't forget your -Concatenate and your -Force switches placed between each switch, and then we go to the server that's being managed. Again, if they're on the same subnet or the workgroup server's network connection is set to Private, we don't need to do anything. If that's not the case, we're going to be changing that rule in the firewall to explicitly allow that computer to connect remotely.

WINRM Windows Remote Management

5:02-6:11

In the subject of remote management, besides working with server pools, Microsoft also includes the WINRM command, and that stands for Windows Remote Management. This is a really cool command. It's used in a number of places. I absolutely adore this. Basically on the server being managed, you open an elevated command prompt and type winrm quickconfig. It's going to give you a couple of prompts, you just hit Yes to both of them. What that does is tell the server that's being managed, it's okay to accept remote commands from somebody else. As long as they're in the same domain, should be no problem.

On the managing server, the way you send the command over is this: you type winrs -r:, and then you put the name of the computer that will be receiving the command, wherever you ran this winrm quickconfig. You put a space, and then you put the <Command>. These are only command line commands. For example, I could do a winrs -r:, the remote computer being dc1, and then run ipconfig. When I get the ipconfig results back, it will be the ipconfig for DC1, regardless of what computer I'm sitting at. You can even do this with an IP address instead of the computer name.

Summary

6:12-6:40

That's the scoop on remote role management. You can go into Server Manager, add in all the servers you intend to manage into a server pool, and then you can install roles and features on anything in the pool without having to remote desktop into the server. We also talked a little bit about winrm, which lets me send over command line commands to a remote computer. The fantastic thing about that is, they won't even know you're sending a command. You could be running an ipconfig, you could be formatting their hard drive, but they've no idea that you're working remotely on their computer. I think that's pretty cool.

2.5.2 Configuring Servers for Remote Management

Configuring Servers for Remote Management
0:00-0:19
In this video, we're going to look at how to configure servers for remote management. We'll be using three servers. DC1 will be managing two other servers. Member2 is a member of the same domain as DC1 northsim.com. Workgroup1 is a member of a WORKGROUP.

Servers on the Same Domain
0:20-0:48
If the server you're going to be managing remotely is a member of the same domain as the managing server, it's very easy. I'm going to click on All Servers. All I need to do is add the server to the server pool. I'll right-click and click Add Servers. There are three tabs. Again, if the server is in a member of the same domain, I can use Active Directory to search for the server.

DNS Tab
0:49-0:54
If not, I can use the DNS tab to add the server by name or IP address.

Import Tab
0:55-1:20
If I wanted to, I could create a text file with an individual name or IP address on each line, and then use the Import tab to import my text file. We'll just use Active Directory tab to add in member2. Now MEMBER2 is a member of the server pool and I can manage it.

Advanced Server Configuration
1:21-1:34
If the server being managed is a member of a workgroup, and the managing server is a member of the domain, or vice versa, then I've got to go through some extra steps to make it possible to remotely manage the server.

Trusted Hosts List
1:35-2:14
On the managing server, I need to add the server that is going to be managed to the Trusted Hosts List. That's done with PowerShell.
I'll set-item wsman -- probably stands for Windows Server Management. On the local host in the client Trusted Hosts List, I'm adding workgroup1. -Concatenate makes sure it gets added to the list instead of redefining the entire list, and -force will make sure that the command goes through. Now, DC1 is set up to manage workgroup1. We need to go over to workgroup1 to set it up to be managed.

Servers are Members of the Same Subnet
2:15-2:19
If the two servers are a member of the same subnet, or if the network adapter is set to private on the server that's being managed, there's no extra step that I need to do.

Network Adapter Set to Private
2:20-2:26

Windows Firewall
2:27-3:24
If not, I would need to adjust the Windows firewall. We'll just take a quick look at what we would do to the firewall, even though in our case we actually don't need to make the change.
I'm going to right-click Windows Firewall with Advanced Security and Run as administrator. I need to adjust an Inbound Rule. The actual rule I need to adjust is Windows Remote Management (HTTP-In). In the Properties of my rule, on the Remote Computers tab, I would add the name of the managing server to the Authorized Computers list. Again, we don't actually need to do that, because they're on the same subnet.
If I'm going to provide the administrator username and password, then at this point I would be done. Let's assume that we don't want to do that.

Add an Account to Manage the Server
3:25-3:58
We're going to create a specific account that will be used to manage this server. Since this is a stand-alone server in a workgroup, I need to create that account in the local SAM, which is done with computer management.
I'm going to create a new user called DC1Admin. I'll add my new account to the Administrator's group. Since my DC1Admin account is a member of Administrators, it will be able to administer the server.

Exempt Remote Users from the UAC

3:59-4:59

However, it would need to be subject to the UAC-User Account Control. Since I'll be coming in remotely, I can't acknowledge the User Account Control prompts, so we need to add a key to the registry that will exempt remote users from the UAC.

We'll add that key to the registry using PowerShell. I'll have a new-itemproperty. The -Name of the new item is LocalAccountTokenFilterPolicy. The -path of the registry is HKLM: (HK Local Machine). Don't forget the colon. \SOFTWARE\Microsoft\Windows\ CurrentVersion\Policies\System. The -propertytype is a Dword and the -value will be 1.

Now, workgroup1 is set up to be remotely managed by DC1. I'm going to add workgroup1 to the server pool; because it's not a member of Active Directory, I need to use the DNS tab.

Adding to the Server Pool

5:00-6:06

You can see, right away, I get an error. The reason is, I haven't provided credentials to manage that server. Whether the server is in the domain or not, if you need to use different credentials to manage the server, you need to go ahead and right-click that server and choose Manage As, then provide the correct credentials.

I'll tell it to remember my credentials so I don't have to do this again. Now I have added Workgroup1 to my server pool and I can do whatever I need to remotely directly from DC1. If I need to use Manage As, it doesn't work well with Remote Desktop Services or the IP Address Management, and in that case, I would have to log into DC1 using the correct account.

That's how you configure servers for remote management.

2.5.3 Deploying Roles on Remote Servers

Deploying Roles on Remote Servers
0:00-0:05
In this video, we're going to look at deploying roles on remote servers. First, you need to have added the servers to the server pool.

Add Servers to the Server Pool
0:06-0:18
If you've added your servers to the server pool, you have them configured for remote management.
It's very easy to remotely deploy a role to those servers.

Add Roles and Features
0:19-0:14

Remotely Deploy a Role
0:15-1:16
I'll simply add roles and features. Then I select the computer from the pool that's going to receive the role. We'll install a role on Member2.northsim.com.
I can choose any role or feature. Before I select my role or feature, let me just take a moment, and take you over to Member2, so you can see that right now, the feature or role is not installed. Right now it's clear, the only role that's installed is File and Storage Services. When I'm done on DC1, we'll come back and see that that's changed. So to Member2, let's go ahead and deploy the DNS Server role.
Now that my installation has succeeded, we should be able to see the role over on Member2.

Confirming Installation
1:17-1:26
After refreshing Server Manager, you can see that DNS has been installed on Member2 by being remotely managed on DC1.
That's how you deploy roles on remote servers.

2.5.4 Offline Images

Offline Images
0:00-0:04
We're going to talk about deploying roles on VHDs and offline images.

Virtual Hard Drive: VHD
0:05-1:12
VHDs are virtual hard drives, so this would be a VHD that you have installed Windows Server 2012 on, but it's not running in a virtual machine. It's just a VHD file just sitting there on the hard drive. We have a couple of requirements. VHD has to have Windows Server 2012 installed in it. It won't work with 2008 R2 or 2008. It's got to be 2012. The VHD cannot have more than one system volume or partition, so it shouldn't have a couple of partitions in the VHD file, because there's no way to specify C: drive versus D: drive.
The network shared folder where the VHD file is stored must be granted access rights. The important thing is, it has to be to the computer account of the server that you're using to mount it. It can't be a user account. It's got to be the computer account. When we go through the demo, you're going to see it written right in the wizard. If this is on a share, that computer account needs Read/Write access to the share. If it's not a share, it's just a local folder, you can really just give Full Control access on the Security tab in the Properties dialog box. It's Full Control NTFS permission.

Using Server Manager
1:13-1:29
When you're going through Server Manager, as we'll see in the demo, first, we're going to click on the destination server page, and we're going to select an offline VHD. It's going to be a radio button. We'll click on the server where the VHD is going to be mounted, and then we'll type in the path where it's located, whether it's a share or a local folder. If you want to do this using PowerShell, it would be install a Windows feature.

Using Powershell
1:30-1:59
You'd have to use -Name and give the name of the feature. Then there's this -VHD switch, and in this spot, you'd put the path to where that VHD is located. If it's share, it's \\blahblahblah, or C:\blahblahblah. If it's on a remote computer, you can put the name of the computer, and in this spot, you'd put the computer name. Then, if you want to be really slick, you can add a -Restart, because a lot of the features need a restart in order for that feature to be installed.

Microsoft Imaging Format: .wim
2:00-3:38
Starting with Windows Server 2008, Microsoft came up with an imaging format. Their style of image is kept in a type of file called a .wim. The sort of sales pitch for this, which I actually agree with, is they say, "It's really great. You can keep multiple images in the same file." It uses single file storage, so if there's a DLL that's used in multiple images, it'll only store the DLL once. It's non-destructive, so when you apply it to a hard drive, it doesn't erase anything that's already on that hard drive, and it doesn't require the disk size to be the same. It's actually used when Windows Server installs, starting from Windows 2008. On the Windows DVD there's a file called, Install.wim that has all the images for all the different versions of Server.
On the Server 2012 DVD, there's an install .wim that's got Foundation. It's got Standard. It's got Data Center. When you put in the product code, it hits an index in that image file and tells it which version to install. What this topic is talking about is, you have some kind of a .wim file that you've made, because you're using Microsoft technology to image new computers. Now you want to add a role, or remove a role, from inside that wim file without having to rebuild the reference computer. I don't have to install the reference image, get the computer up and running, and then change it. I can actually do it directly to the file.
The other neat thing is, the technology we use to do this can also be used on a running computer as well, because from the computer's standpoint, a server that's up and running is just an image that happens to be online. They've really streamlined everything to work together.

Offline Image Role Deployment
3:39-5:03
The first thing you're going to need to do is copy the wim image to the computer. All of my examples assume that we took an image called install.wim and we put it in the c: est folder. You're going to need to retrieve the name or the index number of the image. Like I said, when you're installing 2012, it knows by the product key. You're going to need to go ahead and run a command.
Any time you're working with images, the command is always DISM. If I were in an exam and I got a question on working with images, I wasn't really sure, best bet, at least pick something that starts with DISM. We would get the

image info from an image file. The image file we're getting it from is the C: est\install.wim. That will get back a list of the names and index numbers of the image. You can use either one.
Then we've got to mount it. That's very important. As a generic rule for this test and every other test you're ever going to take, DISM to modify images. Always mount them first. We're going to mount an image. The image file is sitting in c: est. It's called install.wim. The name of the image inside that file, maybe it came back Base Windows Image, or you could say, /index and give a number.
Then we need a mount directory, a directory where we're going to mount it. It's got to be an empty folder on an NTFS partition. I'm going to pretend I made a C:\mount folder. It's empty. I use that in my command, and what it's going to do is put all the files and folders inside that image into that folder.

Find Available Features in an Image

5:04-5:45
We want to find out what features are available, and you'll see a pattern. We're going to go with Running Server Offline Image, a running server, DISM, because it's running, it's online. I just want to Get-Features. It's going to be a big long list of every possible feature that you can turn on and off. If it's an offline image that's been mounted, you've got to say that it's an image, where it's been mounted, then it's the same, Get-Features. If you're not sure about a feature, you can get information about a specific feature. This is online, but it could be offline with the image switch. You just Get-FeatureInfo, FeatureName, and here I just picked bitlocker as one of the features. Again, online would mean online; not online, just make sure you did the image and where it's been mounted.

Enable Features in an Image

5:46-6:05
It's got to be mounted.
To enable features in that image, if we want the feature and sub-features, we need to use the /all switch. Sometimes features are just a feature. Sometimes they have a little arrow you can open up and there'll be sub-features. In a running server, it's DISM. There's our online, Enable-Feature, FeatureName:bitlocker, and if I wanted sub-features, I could add /All.

Patterns in the Commands

6:06-6:48
Hopefully, you're starting to see a pattern in all of these commands. I just want to take a minute to make one comment. In most of the commands, the switches are obvious. In a test situation, if you have to guess, look for the command that looks like that's what it does. You also might want to notice that with DISM, the switches seem a lot like PowerShell commands, verb and then a noun, so that's another trend that you'll see in these commands. If you memorize the trends as opposed to the actual commands, you should be able to pick out the command in a multiple-choice test.
There's our running server. Offline image is going to be exactly the same thing. I've put where the image is. I'm going to Enable-Feature, just like I did for online. Give the FeatureName, all if I want the sub-features.

Status of a Feature

6:49-7:05
Now, if you're curious as to whether your command has gone off okay, you can get the status about that feature by doing the Get-FeatureInfo. I've done it with the online. You can add it to the image. If the status says, Enable Pending, that means the server must be rebooted in order to enable that feature.

Disable Features

7:06-7:30
In a running server, there's our DISM. There's our online. It's just Disable-Feature, instead of Enable-Feature. For the offline image that's been mounted, there's our image with the mount directory, Disable-Feature/FeatureName. You want to get the status to make sure it got disabled; there's your Get-FeatureInfo command that we've seen before. Here, if the status is Disable Pending, then the server must be rebooted to make sure you disabled the feature.

Enable/Disable Using an Answer File

7:31-8:53
With DISM, you can actually work with answer files. Answer files basically give answers. They're usually used during unattended install. You create answer files using the Windows Automated Installation Kit, also known as the WAIK. It's a free download from Microsoft. You create your answer file using a utility named Windows SIM. It will basically walk you right through making an answer file. The answer files get saved with a .xml switch. In the running server, there's our DISM /online. I would Apply-Unattend and the name of the answer file.
This would be some type of feature where it's not easy to get the feature name, or it's a sub sub-feature. For whatever reason, you don't want to mess around with the DISM command line, or maybe you're going to do this again and again and again. You go in and create the answer file and you're just adding maybe one option to the answer file that says, This feature is on. This feature is off. Then, we use DISM to apply that answer file. Historically,

answer files were just used when you were installing Windows to answer any questions during the install. Now, they can actually be used to re-answer different installation options, that way turning them on or turning them off.
In an offline image that's been mounted, this is still just the same exact way that we've been saying it. You put where it's mounted, applying the unattend wherever the unattend lives.

Remove Features from an Image

8:54-9:38
Removing features from an image, there's another section in this course about Features on Demand. This is where we remove the actual files that support the features in order to conserve resources.
If I want to do that to an image in an online server, I can use Dism /online. It's the same disable feature. The key is, you're going to add a /Remove, which is going to completely get rid of the files that support this feature. Offline image, absolutely the same thing. There's my image command. Again, I'm adding the /Remove, all the way at the end. To check that it actually got removed, when I do, I want to Get-FeatureInfo. The status should be Disabled With Payload Removed. They refer to those actual files as the payload. In this case, the payload is gone. The files are completely removed.

Restore Removed Features in an Image

9:39-10:33
Once you remove them, you might want to get them back. DISM is going to use Windows Updates for the source files, unless you add the /LimitAccess switch. Here's a command that would, in a running server or online, enable the feature FeatureName. The only thing different is that we're adding where these files live, the source. I'm assuming Z: is some type of DVD drive. They live in the sources\SxS folder on the Windows Server 2012 DVD. I've added /LimitAccess just so it won't talk to Windows Updates. You can add multiple sources, by the way. You can have as many /source switches as you want.
In the offline image, exactly the same thing. We give where the image is, and then we give our source. If you want to get the status to make sure that that feature that has been removed has been added back in, if the status is Enable Pending, then you've got to reboot it in order to get those files back on and enable the feature.

Saving Changes in Offline Images

10:34-11:16
The very last thing to talk about with the offline images; we made a point of saying we have to mount the image in order to work with it. After you make the changes, you've got to unmount the image. You've got to kind of pack it back up into its wim file and put it to bed. If you want to save the changes that you've made, when you unmount the image, specify where it has been mounted. This is the important part. You would add the /commit switch. If you're not sure or you think you've messed something up, you could actually unmount it, not save the changes, using the /discard switch. As a point of note, if you just run this command and don't say, /commit, and you don't say, /discard, it's just going to give you an error and say, "Hey, you've got to make up your mind. Do you want to keep the changes or do you want to get rid of them?"

Summary

11:17-11:45
That's a whole lot of commands about servicing VHDs and offline images, wim files. But I wanted you to at least have the commands and see some of the trends. Again, when you're working on commands like this, I don't try to memorize all these things, neither for a test nor for real life. In real life, I can do DISM /?, and find out everything I need to know about that command. In the test, I'm going to know that if I'm working with images, I use the DISM command, and I'm going to pick the command that looks like it's the obvious answer to that question.

2.5.5 Deploying Features to Offline Images

Deploying Features to Offline Images
0:00-0:11
In this video, we're going to take a look at deploying features to offline images. We will cover deploying features to VHDs, and also to Windows.wim images as well.

Deploying a Feature to a VHD
0:12-0:29
First, we'll take a look at deploying a feature to a VHD. You have to have a VHD that's running Windows Server 2012, and it can only have one system partition. I've created a VHD like that that's ready to go. We'll go ahead and click "Add roles and features". First, we select the computer where the VHD is going to be mounted.

Mount VHD to Selected Computer
0:30-0:59
I also need to specify that I'm going to be managing a virtual hard disk.
If I do that, I can go ahead and browse, and find the VHD. If it was on a different server, I could also put the UNC path name of the VHD. In my case, mine is Local. It's in the C: Drive, and I've named it Server2012 VHD. Once Server Manager connects to the VHD, I add roles just like I would if the server were online and running.

Adding Roles (Example: Adding DNS)
1:00-1:31
Let's add DNS. It can't determine what the IP address is in the VHD, so it's giving me a warning. I'm not going to worry about that. Now, DNS has been installed on the server in that VHD, even though the server isn't even running. Now, we're going to look at deploying a feature or role to an offline image using the Windows image format.

Using Windows Image Format to Deploy Features
1:32-2:09
I have a Windows Image in the C: test folder, and it's named install.wim. I've just copied the one from the Windows Server 2012 DVD. It's in the Sources folder on there, and it's the one that's actually used to install Windows, but we can work with it directly as an image as well. This could also be a custom image that I had created using ImageX, or whatever I would use to create a WDS. I can work with any .wim image the same way that we're going to do it just now.

Mounting the Image
2:10-3:43
Before I can edit that image, I need to mount it. To mount it, I need an empty folder on an NTFS partition. We'll make a new folder named mount. You can see that my new folder is empty. To work with a WIM image, we need a command prompt. I'm going to go ahead and run it as administrator. Right-click, Run as administrator. I can't change the image until it's been mounted. In order to mount it, I need the name, or index number, of the image. This tells me all the images that are inside of that install.wim, and gives me the index number. When I mount the image, I can use either the index number or the name, whichever is more convenient. We'll go ahead and work with the data center image, which is index number 4.
My next step is to mount the image. If I wanted to work with the name, instead of index, I would just use name, and then put the name of the image in quotes. You can see now that the image has been mounted. My mount directory has been populated with the files in that image. If all I needed to do was to copy a file to the hard drive, I could just put it right in this folder. In order to make changes in the operating system, I've still got to continue to use DISM.

get-features
3:44-4:01
One thing I can do is get a list of all the roles and features that I could possibly manage using DISM.
You can see that the DISM image, I specify where the image is mounted. Get-features gives me a list of all the features that I can control using DISM.

get-featureinfo
4:02-4:21
If I want to get more information about a feature, I can do that as well. I'm going to get-featureinfo about the DNS server role. If there were subcomponents, or anything else I needed to do, it would be listed here.

Enable/Disable Feature

4:22-4:59

Now that I know the name of my feature, I can use DISM to turn it on. DNS has been turned on in that image, even though the computer is not even running. If I needed to turn it off, instead of enable feature, I would disable feature.

I could also use an unintended install file and apply that to my image as well. Now, I'm done. I've turned on what I needed to turn on.

Unmount the Image

5:00-5:55

The last thing I need to do is unmount the image, and basically pack up all the new files back into that same file. When you unmount the image, you have to choose whether to keep the changes or get rid of the changes. If I just run this command, I'll get an error. The reason being, I either need to commit the changes, save them, or if I'm afraid I've made a mistake, I could discard the changes as well. We'll go ahead and save them.

You'll see as it unmounts, the files in the mount directory will disappear. The unmounting has come back with an error, but that's actually not true. It might be because I have my mount folder open, but it is dismounted. We can see that the folder is empty, so we're all set. DISM can also be used to make changes to a computer that's actually running.

Using DISM to Manage Computers that are Running

5:56-6:25

For example, if we look at my computer here, it's not running DHCP, but I can turn it on using DISM. Now, if I go back into Server Manager, and I refresh, you will be able to see that this enabled-feature has installed the DHCPServer on this computer. There's DHCP, right there.

Summary

6:26-7:10

We've looked at deploying features and roles to offline images. We saw that if our image is in a VHD, we can just use Server Manager, like we would any remote computer whether it's running or not. If the offline image is in a WIM file, we've got to use DISM, and we can even use DISM to manage the computer while it's running. If you just know that we work with offline images using DISM, and that most of the switches are pretty obvious, that should be good enough. When it comes to exams, you've got to pick your battles, and I don't think memorizing long DISM commands is going to be as beneficial as just knowing the commands start with DISM. In looking for the command that makes sense. That's how we deploy features to offline images.

2.5.6 WinRM

WinRM
0:00-0:36
In this video, we're going to take a look at using WinRM -- Windows Remote Management -- to manage computers. Since we're going to be managing some of the servers remotely, we need at least two. DC1 is going to be the managing computer. Member2 is going to be the computer that we will be managing. If they're all Windows Server 2012 and they're in a domain, they may already be set up for remote management. We're still going to take a look at the command that we get a computer enabled for remote management, just in case it's an older computer or it's not in the domain. Pretty easy to do.

Open a Command Prompt
0:37-1:07
First, we want to open the command prompt, and I'm going to make sure that I right click it and choose Run as administrator. The command to get the computer setup to be managed is winrm quickconfig or you can also use qc. Because it's Server 2012, they're in a domain. It comes back, it's already running on this machine, it's already set up for remote management, but if this were Windows Server 2008 R2, we'd get a couple prompts -- one to turn on remote management, the other to open up a port in the firewall.

Looking at the Features of a Computer
1:08-2:39
While we're here, let's take a look at some of the features of this computer so that we'll know when we see them again.
We can see that the host name is in fact Member2. Member2 right now has an IP address of 192.168.1.51, and here are a list of the folders in the root of the C: drive. We'll also go into Disk Management and take a look at the hard drives. You can see I have my Disk 0, there's my system partition, which is used to boot the computer, and there's my C: drive, which is the boot partition where Windows is installed. I have another disk that's offline. We can bring it online, and it's got roughly 127 GB. Right now, all of that space is unallocated. We'll minimize computer management.
Let's go over to our managing computer DC1. We'll open up the command prompt as administrator and take a look at some of the settings on this computer. Its host name is DC1 and its IP address is 192.168.1.40. Here's where the cool part comes in.

Sending a Command from Managing Computer to Remote Computer
2:40-3:32
We can send any command line command from DC1 to Member2. When I put in the winrs, the -r: specifies the name of the remote computer. I can specify it by name, or I could use the IP address. I'll leave a space, and then I put the command line command I want to send over to that computer, and it comes back with that as Member2, which is exactly what we saw when we were local to Member2. We do an ipconfig and we see the IP address of Member2.
If you really want to get slick -- I like this -- if you put cmd as your command, it's going to throw the whole command prompt over to the context of Member2. You can see up at the top it says the command prompt is using winrs to Member2. Everything I do is done on Member2. Just to be fun, we'll go ahead and mess around with the hard drives.

Example: Modifying the Hard Drive
3:33-4:36
There are the two hard drives that we saw when we were setup over on Member2. Disk 1 was that second drive, and we can go ahead and work with that. What I've done is gone ahead and make a partition. The partition is primary partition, roughly close to 20 GB. Notice with this part, everything is singular. It's not list disks, it's list disk list partition.
We'll do a quick format of our partition, and now that I formatted the volume, I'll give it a drive letter. Now that I partitioned the hard drive, the last thing I'm going to do is go into the root of the C: drive and make a new folder. If we want to get out of our winrs command prompt, we can do an exit that puts me back in my regular command prompt, and now I'm done.

Getting out of WinRS Command Prompt
4:37-4:44

Results from Example
4:45-5:27

Let's go take a look at what happened to Member2. If I go in and I look at the C: drive, there's my remotely made folder, there is the E: drive that we created -- about 20 GB -- and if we take a look in Disk Management, we can see that.

Notice, nothing changed on the screen of Member2 until we came back to take a look around. Because I'm working with the command prompt, anybody who is logged in locally isn't going to see anything, and that's how you can use WinRM to remotely manage another computer -- and make sure you understand it's only command line commands I can send over.

2.5.7 Remote Management Facts

Remote role deployment in Windows Server 2012 allows you to install and manage roles on multiple remote servers. Keep in mind the following about remote role deployment:

- Server Manager can add roles and features only on remote servers running Windows Server 2012.
- Server Manager can manage remote servers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.
- Server Manager can manage up to 100 remote servers.
- Limiting the amount of event data collected by Server Manager increases server efficiency when managing large numbers of servers.
- Server Manager can receive only online or offline status from servers running Windows Server 2003.

To deploy roles to remote servers:

- Use Server Manager to create a *server pool* of the remote servers to manage.
- Use **Manage As** to change credentials when required to access a remote server. The following roles do not support **Manage As**:

  Remote Desktop Services (RDS)
  IP Address Management (IPAM)

  When using these roles, you have to log on with the account that has rights to those servers.

The following table identifies remote role deployment tasks:

Task	Description
Add domain servers to the server pool	To add domain servers to the server pool, right-click **All Servers** in Server Manager: • Use the **Active Directory** tab to add servers from the domain to the server pool. • Use the **DNS** tab to add servers by name or by IP address. • Use the **Import** tab to add multiple servers using a text file. List the name or the IP address of each server on a separate line.
Manage non-domain servers	To install and manage roles on a workgroup server (a server not in the domain): 1. Add the workgroup server to the **TrustedHosts** list using the **Set-Item** command. For example, to add a server named **Kiosk1**, enter: **Set-Item wsman:\localhost\Client\TrustedHosts <KIOSK1> -Concatenate -Force** 2. Configure the workgroup server firewall:

	○ If the workgroup server is on the same subnet as the managing server, or if the workgroup server's network connection is **Private**, no change is necessary. ○ When the server is on a different subnet and the network connection is not **Private**, change the inbound firewall rule **Windows Remote Management (HTTP-In)** to include the name of the managing server. 3. Enter the credentials for the workgroup server. The options you have are: ○ Use the Local Administrator credentials. ○ Create an account with administrative privilege. When using this option, you must override the UAC by creating a new registry key: **New-ItemProperty -Name LocalAccountTokenFilterPolicy -path HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System -propertyType DWord -value 1** To install and manage roles on servers in the domain from a workgroup server not in the domain, use this same process. In Step 1, you add the name of the domain servers to the workgroup server.
Deploy a role on a remote server	To deploy a role to a remote server: • Select the installation type. • Select the server to receive the role from the server pool.
Deploy a role to a VHD file	To deploy a role to a VHD file: • The target VHD file must meet the following requirements: The target VHD file must have Windows Server 2012 installed. The VHD cannot have more than one system volume or partition. • The computer or local system account of the server mounting the VHD requires: **Read/Write** access on the Share where the VHD is located. **Full Control** access if stored locally. • In the Add Roles and Features Wizard, select a virtual hard disk as the destination. Select the server on which to mount the VHD. Enter the path to the VHD file. • You can use the PowerShell command: **Install-WindowsFeature**. For example, to install the DNS role in a file located in file **F:\offline** on a computer named **dc1**, use the following command: **Install-WindowsFeature -Name DNS -VHD F:\offline -ComputerName dc1 -Restart**

The Windows Remote Management WINRM command allows a server to accept remote commands from another server in the same domain. To use WINRM:

- On the managed server use the **winrm quickconfig** command.
- On the managing server, you send commands in the following format: **winrs -r:*remoteserver command.*** For example, to send the **ipconfig** command to a remote server named **dc1**, enter

  winrs -r:dc1 ipconfig

You can use Deployment Image Servicing and Management (**dism**) commands to:

- Deploy a role to an offline image
- Find features in an image
- Enable features in an image
- Get the status of a feature
- Disable features in an image
- Remove features from an image
- Restore removed features in an image
- Unmount an offline image

You can also use answer files with the **dism** command.

- Answer files are typically used for an unattended install.
- To use an answer file, install the Windows Automated Installation Kit (WAIK).
- To create an answer file, use the Windows SIM utility.
- Answer files are saved as .xml files.
- On a running server, run dism and identify the answer file. For example, to use an answer file named myunattend.xml, the **dism** command is:
  dism /online /Apply-Unattend:C:\answerfiles\myunattend.xml

2.6 NIC Teaming

As you study this section, answer the following questions:

- What are the benefits of using NIC Teaming?
- How does switch-dependent mode differ from switch-independent mode?
- How does static teaming differ from dynamic teaming?
- What is the function of the Link Aggregation Control Protocol (LACP)?
- How many standby NICs are there in an active/passive configuration?
- What are the restrictions of using NIC Teaming with Hyper-V?

After finishing this section, you should be able to complete the following task:

- Configure NIC Teaming.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  Configure Server Services
  - Configure NIC Teaming

This section covers the following 70-410 exam objective:

- 102 Configure servers.
  This objective may include but is not limited to:
  Configure NIC teaming

2.6.1 NIC Teaming

NIC Teaming
0:00-0:08
In this video, we're going to take a look at NIC teaming. NIC teaming allows me to combine two or more network adapters together, where they work as a team.
Reasons for NIC Teaming
0:09-0:11
I might want to do this for two different reasons.
More Bandwidth
0:12-0:33
One reason is, maybe I need more bandwidth. When you create the team, if both members of the team are active, then you'll have the aggregate bandwidth available to the clients. If I have two network cards, they're attached to 100 megabits per second ports in the switch; I create a team. Essentially, that team has 200 megabits per second available to them going through that switch.
Fault Tolerance
0:34-0:57
The other reason might be for fault tolerance. Maybe I have two network cards on my server, not because I have too much traffic going in and out, but because I'm afraid that if one network card fails, communication to the server will be cutoff. I can use NIC teaming and bind these two network cards together, where if one fails, the other one will either keep working, or if it hasn't been working, it will take over for the primary.
Load Balancing and Failover (LBFO)
0:58-1:07
NIC teaming is also known as Load Balancing and Failover, LBFO. You should be aware of that acronym so anytime you see LBFO, think NIC teaming.
NIC Teaming
1:08-1:24
What it does is allow multiple network adapters on the computer to be placed into a team. There's my bandwidth aggregation. I'm looking for total bandwidth of the two network cards or I'm looking for fault tolerance, so that if one dies, we'll have failover to the back up. If they're both working as a team, the second one will continue to work.
NIC Team Appears as a Single Adaptor
1:25-2:12
When I go through and I create a NIC team, it appears as a single adapter to the operating system. You definitely need to know that.
If I'm going in to Hyper-V and creating switches, each external switch can use one team, because Hyper-V is going to see this as one network card. Hyper-V is not able to create networks based on the individual adapters in a team. If you have one team of two adapters and the boss comes in and says, "I really need two external virtual network switches in Hyper-V", you're going to have to break the team, and here's how you would do it from PowerShell. Remove-Net, there's my LbfoTeam, and then I would specify the team. I can do this in Server Manager, or if I love PowerShell, I can use the New-NetSwitchTeam command, Requirements.
Requirements.
2:13-2:14
Presence of at Least One Network Adapter
2:15-2:33
Networking teaming requires the presence of at least one Ethernet adapter. You might say, why do I want to do a team with one adapter? Maybe I'm trying to separate traffic using VLAN. Maybe I have one physical adapter, but I go into my virtual machine, I give that two virtual adapters, make them into a team, but each member uses a different VLAN.
Presence of at Least Two Network Adapters
2:34-2:53
If you're looking for that fault protection through failover, then you need at least two Ethernet network adapters. That really is fault protection or fault tolerance, minimum of two of everything. Windows Server 2012 supports up to 32 network adapters in a team. In Hyper-V, they only support two. In the host server, you can have 32.
Two Scenarios
2:54-2:53

Switch Dependent Mode

2:54-3:53

When I'm setting up my teams, I have two scenarios. I might have a host with all adapters connected to the same switch, and this is called Switch Dependent Mode. But they're all connected to the same switch, there is no configuration protocol, which is called static or generic teaming, or we can have a mode that uses this LACP to coordinate between the host and the switch, in which case, we are going to have to have some configuration on the switch.

If they're all connected to the same switch, we're probably looking for bandwidth aggregation, so we're going to combine these to get the maximum bandwidth. All the packets associated with a single TCP stream or single conversation would be handled on one of the NICs. The reason we want that is this: if we had one conversation coming in to either member of the team, it's possible we could get out of order packets. If the conversation is packets one through ten, it's possible that one adapter might process packet 5 before the other one is done with 1 through 4 and that would be a problem. In this scenario, the teams are usually Active/Active.

I also could have a situation where I have a host with each adapter connected to a different switch.

Switch Independent Mode

3:54-4:40

This is where I'm looking for high availability. Not only am I worried about the network card going down, I'm worried about the switch going down-- it's called Switch Independent Mode, and the switches are not aware that different interfaces on the server make up a team. It could be Active/Active or Active/Passive. If you have an Active/Passive, there's only going to be one standby network card per team.

We'll see in there Generic or Static teaming, which means I've got to configure both the switch and the host. I have an option for LACP which enables the automatic creation of a team, but I'm going to have to enable that on the port. It's a little bit more work, and you have to have a switch that supports it. I would say, if you see LACP, just think, it makes it a little bit more dynamic, has to be configured on the switch, you should be good to go.

Traffic Distribution Algorithms

4:41-5:20

The last thing we'll talk about in terms of concepts are Traffic Distribution Algorithms. If I choose Hyper-V switch port, then the VLAN's MAC address or the port it's connected to on the Hyper-V switch is the basis for dividing traffic. If not, they can do this address hashing, which makes a hash based on components of the packet, and then assigns packets that have that hash value to one of the available adapters. Remember, we're trying to keep things within the same stream on one adapter. It can create the hash out of any one of these things: source and destination MAC addresses, source and destination IP addresses, TCP ports and IP addresses. This is your best option, but it cannot be with the IPSec. If you need an IPSec, you can't select this.

Teaming Mode

5:21-5:24

When we go through the actual wizard, the first thing we're going to be asked for is the teaming mode.

Static Teaming

5:25-5:31

Static teaming does require configuration on the switch and the computer. We identify which links form the team, and it's switch dependent.

Switch Independent

5:32-5:39

Switch independent lets me distribute the NICs across numerous network switches for the ultimate in fault tolerance, and then I've got this LACP which gives me Link Aggregation, and I can expand or reduce the team; that's also switch independent, except the port has to support it.

LACP

5:40-5:51

I don't have to configure the switch and tell it these things are attached, but I just have to configure the port.

Load Balancing Mode

5:52-5:53

Next thing I'll be asked for is the Load Balancing Mode.

Address Hash

5:54-6:06

Here's our address hash again, which lets me load down to the network traffic. This is your best choice for most configurations, and the great thing about it is, it does not disrupt communication between the VM and the network if one of the network cards fail.

Hyper-V Port

6:07-6:41

I can also choose Hyper-V port, where I can load balance it by virtual machine, instead of trying to say, "well I want the TCP stream on one NIC," I can say "I want this virtual machine on one NIC," but that virtual machine is only going to transact over one of the NICs in the team. If you have multiple virtual network cards in your VM and it's teamed inside the guest operating system, that's the scenario where you would choose Hyper-V Port. If you're not doing a team inside of the guest operating system, go with Address Hash.

If I got a failure of the network card that the VM is using, that could potentially disrupt communication between the VM and the network.

Summary

6:42-6:58

NIC teaming allows me to connect at least one--but usually two or more--network adapters and create a team, and I'm either looking to aggregate bandwidth, or I'm looking for failover, where if one adapter stops, the other will take over. That's what you need to know about NIC teaming.

2.6.2 Configuring NIC Teaming

Configuring NIC Teaming
0:00-0:02
In this video, we're going to take a look at NIC teaming.
Creating a NIC Team
0:03-0:36
When you set up a NIC team, you do it inside of Server Manager. I'm just going to go over to Local Server and you can see that NIC teaming is Disabled. If I click on it, it brings up the NIC teaming dialog box, and the first thing that I want to do is go ahead, click on Tasks, make a New Team.
Give our team a name; we select the physical adapters that will be a part of the team. If I were inside of a virtual machine I'd be limited to two. On a physical machine I can have up to 32 working together.
Teaming Mode
0:37-0:40
The first thing that we select is our Teaming mode.
Static Teaming
0:41-0:53
Static Teaming means that they're plugged into the same switch and we're going to coordinate with the switch to let it know that these two adapters are teamed together, and we're going to aggregate the bandwidth.
We're going to get the combined bandwidth of both ports.
Switch Independent
0:54-1:06
Switch Independent means they're plugged into different switches. I can do aggregation or I can just use it for failover. Maybe I just want to make sure that if one network card or one switch dies, the other one is available.
LACP
1:07-1:14
LACP means that the computer will coordinate with the switch and figure out the teaming.
We're going to go Switch Independent.
Load Balancing Mode
1:15-1:31
Now I choose my Load balancing mode. With Hyper-V Port, that means that the computer will track the traffic by virtual machine, so that any particular virtual machine will be attached to only one of these adapters, and use just that one adapter for traffic.
Address Hash
1:32-1:17
Hyper-V Port
1:18-1:49
Address Hash means it's going to make a hash for each TCP stream and make sure that each stream is processed by one of the adapters.
The reason for that is, we're trying to avoid packets arriving out of order and messing up the conversation. I'm going to leave it with Address Hash.
Standby Adaptor
1:50-1:59
I can also set up to one Standby adapter, so I can have them all Active, or I can have up to one backup for the whole team. We're just going to leave them all Active.
Making Changes
2:00-2:12
Now you can see our network team is working. If I need to make any changes, I can go into the Properties and make any changes that I need to make.
Once I create the team, it appears as if it's one adapter to the operating system.
Adaptor
2:13-2:45
Let's go in and take a look. You can see that these two adapters here are working, but if we go in, all the computer really is working with is the team. And let me show you that. If I do an ipconfig; you can see that there's only one adapter as far as the operating system is concerned.

Summary

2:46-3:18

When we enable NIC teaming, we're tying together two or more network cards -- either, because we want to combine the bandwidth of both cards, or we're looking to do failover. The team itself appears as if it's one adapter to the operating system. Another thing to know if you have an adapter that's being used in Hyper-V as an external network -- you cannot add that adapter to the team.

If you want to base your external networks on the team, you've got to create the NIC team first, then create your Hyper-V networks. That's how we setup NIC teaming.

2.6.3 NIC Teaming Facts

NIC Teaming allows two or more network adapters to be combined to work together as a team. NIC teaming can increase bandwidth and provide fault tolerance. When you create a team of active network adapters, the aggregate bandwidth of the network adapters is available to the client. Teamed network adapters can be configured to provide load balancing and fail over (LBFO).

Be aware of the following regarding NIC Teaming:

- Windows Server 2012 supports up to 32 network adapters in a team.
- A NIC Team appears as a single adapter to the server operating system.
- Requirements for NIC Teaming are:
  - At least one Ethernet network adapter. You can use one adapter to separate VLAN traffic.
  - A least two Ethernet adapters are required to provide fail over.
- When using NIC Teaming with Hyper-V:
  - Each external switch can use only one team.
  - An adapter connecting the VM to the network cannot be part of a NIC Team. You must create the NIC Team first and then create the Hyper-V network.
  - Hyper-V is not able to create networks based on individual adapters in a team.
  - Hyper-V supports only two adapters in a team.
- Use PowerShell or Server Manager to manage NIC Teaming. In PowerShell:
  - Use the **New-NetSwitchTeam** cmdlet to create NIC Teams.
  - Use the cmdlet: **Remove-NetLbfoTeam** to break up a NIC Team.

NIC Teaming can be set up in one of two ways:

Mode	Description
Switch-dependent mode	Switch-dependent mode requires all network adapters to be connected to the same switch: • All NICs in the team are connected to the same switch using one of the following methods: *Static* or *generic* teaming requires that the links forming the team be identified on the switch and the computer. *Dynamic* teaming uses the IEEE 802.1ax Link Aggregation Control Protocol (LACP) to identify the links that form the team. Most switches require that LACP be manually enabled on the port. The LACP protocol is also referred to as IEEE 802.3ad. • The bandwidth of the adapters is aggregated. • Traffic distribution should be implemented so that packets associated with a TCP stream are handled by the same network adapter.

	• The teams are usually Active/Active meaning that both network adapters accept traffic.
Switch-independent mode	In switch-independent mode each adapter is connected to a different switch. Switch-independent mode provides fault tolerance. In this mode: • Switches are not aware of the NIC team. • The NIC team can be Active/Active or Active/Passive. In an Active/Passive configuration there is only one standby NIC per team. • NICs in the team can be connected to the switch using either static teaming or dynamic teaming.

The following table describes the *load balancing mode*, also known as *traffic distribution algorithms.*

Method	Description
Hyper-V switch port	The MAC address can be used to divide traffic when virtual machines have independent media access control (MAC) addresses. • The advantage to this method is that the switch balances the traffic based on the MAC address for the virtual machine. • A disadvantage is that the virtual machine is limited to the bandwidth of a single adapter. • Choose Hyper-V switch port if you have multiple virtual network cards in the VM teamed in the guest operating system.
Hashing	The hashing method creates a hash for the packet and sends packets with that hash value to an available network adapter. • Dynamic redistribution of packets based on hash value is known as *smart load balancing* or *adaptive load balancing*. • Hashing ensures that all packets from the same stream are sent to the same network adapter. • Communication between the VM and the network is not interrupted if one of the adapters fails. • The hash is created using one of the following: Source and destination MAC addresses. Source and destination IP addresses.

	Source and destination TCP ports, and source and destination IP addresses. This type of hash cannot be used with IPSec.

2.7 Traditional Storage

As you study this section, answer the following questions:

- What are the advantages of GPT partitioning over MBR partitioning?
- Why is MBR recommended for disk sizes smaller than 2 TB?
- How does a basic disk differ from a dynamic disk?
- Which disk configuration is used with spanned, striped, and mirrored volumes?
- What information does the dynamic disk database contain?
- When would you use a mount point?
- How do spanned disks use disk space?

After finishing this section, you should be able to complete the following tasks:

- Initialize a disk.
- Configure a volume.
- Extend a volume.
- Configure fault tolerant volumes.
- Create a mount point.
- Manage disks and volumes.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  Configure Server Storage
  - Configure Server Volumes
  - Configure Fault Tolerant Volumes
  - Create a Mount Point

This section covers the following 70-410 exam objective:

- 103 Configure local storage.
  This objective may include but is not limited to:
  Configure basic and dynamic disks
  Configure MBR and GPT disks
  Manage volumes

2.7.1 Disk Storage

Disk Storage
0:00-0:06
In this video, we're going to talk about disk storage. We're going to talk about adding disks and storage pools.
Initializing
0:07-0:14
When you first add a disk to the computer, you've got to initialize it, that will show up as needing to be initialized. There's two ways you can initialize it.
MBR
0:15-0:23
MBR is the type of initializing we've used since the beginning, so it's good for backwards compatibility, but the only problem with it is, it only supports disks up to two terabytes.
GPT
0:24-0:39
That might not be a problem for you, but if it is, and you have disks that are two terabytes or larger, you need to go GPT, and that supports disks that are larger than two terabytes.
Here's the only problem: it came in with Windows Server 2008, so it's supported by 2008, 2008 R2, and 2012.
How to choose between MBR and GPT
0:40-1:34
You might be thinking, GPT, that's the new one. Let me go with that. Just a word of warning: it's better to go MBR, unless you actually need GPT. If your disk isn't two terabytes or larger, go MBR, and here's a good reason why. If that disk fails and for any reason you need to send it out to a data recovery expert, a lot of the data recovery plans don't run in Vista or better. They actually run on the older operating systems, like sometimes as far back as Windows 98, Windows XP. So if you choose GPT, you may not be able to use a data recovery expert. If you choose MBR, you should be all set.
That might not be an issue, maybe you'll say, Well, I'll never do data recovery. That's up to you. I always like to keep my options open. Even Microsoft will recommend, if I don't have a disk two terabytes or larger, might as well just go MBR. You can convert between them as long as there are no volumes on the disk. Once you start creating volumes on the disk, then you're pretty well set.
Storage Spaces/ Storage Pools
1:35-2:30
We're going to talk about storage spaces, which is new with Windows Server 2012. Basically, what it does is allow physical disks to be managed in storage pools. For my purposes, storage spaces, storage pools, pretty much mean the same thing. The storage pools allow your storage to be managed dynamically. Disks can be added or removed from the pool as necessary, but in order to add a disk to the storage pool, the disk must be online and unallocated. That you have to know, and it's going come up again in the list of limitations.
As soon as you add a disk and it's initialized, sometimes even before it's initialized, it's going to show up in this existing pool, called the primordial pool. It's not a real storage pool. It's just a place where they stick the disks until they actually go through and put them in our own storage pools, if we're going to use them.
After you add a disk or a physical disk to the storage pool, it's going to disappear from Disk Management. You'll have to create a virtual disk in the pool, and only virtual disks created from the pool will be available.
Hardware Raid and Storage Pools
2:31-3:39
If you're familiar with hardware RAID-- when you have hardware RAID, basically, you have a hardware card or RAID controller that's managing the disks.
Let's say I have RAID 5 card, and I've got seven disks in my RAID 5 array, and maybe, let's say, six of them are active, one's a spare. Traditionally, to the operating system, it looks like I've got one big hard drive. To make the math easy, let's say each of those disks is just one terabyte. It would look to the operating system as if I have a six-terabyte- big hard drive, and then that seventh disk that we're using as a spare- is just going to be hanging out. If one of the other disks dies, we can use the space to repair the array.
What storage pools are really is, they mimic this hardware RAID. If I add the seven disks to my server, I go into Disk Management; I'm going to see seven disks. If I add them to a pool, I can literally go into Disk Management, I wouldn't see any of those disks, but then I go into the storage pool and I create a virtual disk, and that's what shows up in Disk Management, the only difference being with hardware RAID--usually it's a little bit difficult to add disk. I've got to break the array, restart it, depending on what you vendor allows. Here, we can expand that storage pool as much as we need to.

Creating a Storage Pool
3:40-3:41
Our steps for creating a storage pool: the first thing we're going to do is add our physical disks to the storage pool.
Step One: Add Physical Disks to the Storage Pool
3:42-3:48
That will move them out of primordial pool into the new pool that I create.
Allocation
3:49-4:15
My first choice is going to be allocation, and I've got two choices: automatic, which adds the space from the disk into the pool. Hot spares are just that. It's that disk, or the space on that disk is going to be used only if one of the automatic disks fails. It's good to have a hot spare if you can. You have a hot spare, then that means if one of the disks fails, essentially, you can just keep going without a LAN administrator having to be there and putting a new disk in. If I don't have a hot spare, what's going to happen---it depends on what I choose later on in the wizard.
Step Two: Create a Virtual Disk
4:16-4:21
My second step is to create a virtual disk. When we go through the virtual disk wizard, we're going to have some choices.
Storage Layout
4:22-5:15
The first choice we're going to see is storage layout, and these are very important. You really want to know this if you're going to be taking any tests of any kind. A simple storage layout is the same as striping and to spread the data across the disks. The only problem being, if any of those disks die, then that virtual disk is going to go down. With a mirror storage layout, the data will be mirrored across at least two disks, so it requires at least two disks, but because 100 percent of the data is saved on both, I can lose one disk and still keep going, even if I don't have a hot spare.
With the parity disk, that's really the same as RAID 5. RAID 5 requires at least three disks, and the data and the parity information will be spread across those three disks. Parity is information that we can use to rebuild data that's missing, so the RAID 5 up to one disk can fail, and the virtual disk will keep working. If there's a hot spare or I replace the failed disk, it will rebuild the missing data that was on that disk.
Provisioning
5:16-6:57
The next thing we're going to be asked about is provisioning. We've got two types of provisioning. Fixed is what they call it in the operating system, but you'll see documentation from Microsoft that might call it thick. Basically, it uses all the available space in the storage pool. That means however big I make this virtual disk, all that space is going to be subtracted from the pool and used up immediately.
Thin is a little bit more interesting. If I use a thin virtual disk, I can actually make my virtual disk larger than the space available in the pool. Let's say I have five terabytes in the pool, theoretically, I can go through and make a virtual disk that's 10 terabytes. Here's the catch: that virtual disk is going to grow and grow and grow. You need to make sure if it grows beyond five terabytes, you've added more physical disk to the pool, otherwise you've got a math problem. You're going to be very sad.
If you're using fixed disks and the storage layouts are not simple, then there's going to be a little bit of space consumed to set up that storage layout, so it's going to consume more free space than the size specified. If you can use the maximum, that's great, if not, you've got to accommodate for the fact that some of that space is going to be used up to create the storage layout. By default, you can create a virtual disk only if there's sufficient free space to do that. You've got to have at least some of the free space available. If it's thin, it gets a little bit more tricky, but it's got to have enough space to create the virtual disk.
I'd be very, very careful with the thin provisioning. That's going to require a lot more monitoring, and certainly I can see the potential to get yourself in trouble there. Let's say you've got five terabytes now, you make your virtual disk 10 terabytes, but you know you've got those other five terabytes coming in the mail. That gives you an option to account for that. You've got some flexibility there, just keep an eye on it.
Limitations
6:58-6:59
Here's some limitations.

No Boot, System, or Cluster Shared Volumes (CSV)

7:00-7:49

First of all, we cannot have the boot system or any clustered shared volumes on the storage pool. The boot partition is where the operating system files are kept. This is kind of a famous Microsoft oxymoron. This is whatever drive the C:\windows folder is in, that can't be in the storage pool. The system partition is a partition that has the files necessary to boot the operating system. When you Next, Next, Next, Finish, if you just install server that way, it creates a little mini partition, it's usually like 100 mgs, and that's your system partition. That can't be in the storage pool either.

Clustered shared volumes are used to support live migration with Hyper-V and it's Hyper-V, integrated with the clustering service and then you make a clustered shared volume. If you have those, you can't put that in a storage pool either.

Minimum Drive Size: 10GB

7:50-7:52

The individual drives in the pool must be 10 gig or larger.

Only Add Un-formatted/Un-partitioned Drives

7:53-8:22

When you introduce a drive into the storage pool, the contents of that drive being added will be lost. Again, it's very important to know, you want to add only unformatted, un-partitioned drives. You should have unallocated space.

If I gave you a scenario and I said, Here's Drive2. It's got a partition on it, the E: drive. It's formatted with NTFS. I want to add in into my storage pool. How would I do that? The answer would be get rid of that partition, because we want un-formatted, un-partitioned, unallocated space before we add it.

Needs at Least One Drive

8:23-8:32

A simple storage pool has to have at least one drive. If you've only got one drive there's no point making a storage pool, but you can start your pool with one drive knowing you're going to expand later.

Fiber Channel and iSCSI Drives Not Supported for Failover Clusters

8:33-9:08

Fiber-channel and iSCSI drives, which are drives out on a storage area and network, they're network attached storage in a storage area network, are not supported if you're using these for clusters. Finally, virtual disks that are going to be used with the failover cluster that come from a storage pool have to have the NTFS file system. New with Hyper-V 2012, it supports just failing over the VMs, and all they're saying is, if you're going to have a virtual disk or a virtual machine and it's going to be on the virtual disk in the storage pool, then it's got to be NTFS. It gets a little confusing, but it's really not that bad. The operating system will correct you if you do something wrong.

PowerShell Commands

9:09-9:41

Last, we'll take a look at the PowerShell commands. If you wanted to make a new storage pool, the command's pretty simple, New-StoragePool. Add-PhysicalDisk adds a physical disk to an existing storage pool. Once I add my disk to the pool, I have to create virtual disks, so this will make a New-VirtualDisk. Get-StoragePool retrieves information about the storage pool. You can see New- or Add- is always for creating something new. Get- is always for retrieving information. If we had a set command that would be for changing settings on something that already exists.

Summary

9:42-10:10

That's how we manage disks in Server 2012. We've got to initialize them; that's been true for quite a while. I can use MBR, which is the industry standard. If I go two terabytes or larger, I'm going to use GPT. If I want to go through and get maximum flexibility, I can create a storage pool, which essentially lets me manage my drives as if I were using hardware RAID, except I can add a disk to the array much more easily. Once I've got my storage pool, I create virtual disks. Make sure you know the difference between simple, mirror, and parity, and then on those disks, I'll create volumes.

2.7.2 Disk Facts

Directly-attached storage is commonly implemented in the following ways:

- Internal hard disks using an internal bus and IDE, SCSI, or SATA drives
- External USB, eSATA, or FireWire disks
- Rack mounted storage connected to the server through an external connector or external bus

When you add a new disk to a Windows Server 2012 system, you will be prompted to initialize the disk and select a partition style. The partition style identifies how information about disk partitions are stored. The following table identifies the two partition styles of disks:

Partition Style	Description
Master Boot Record (MBR)	Master Boot Record (MBR) is the traditional partitioning method. MBR provides: • Backwards compatibility with Windows operating systems. • Support for up to 4 partitions. • A maximum partition size of 2 TB. With MBR, partitioning information is often stored in hidden locations on the disk or even within the operating system. These locations might vary from vendor to vendor. If you do not have a disk larger than 2 TB, use MBR.
GUID Partition Table (GPT)	GUID Partition Table (GPT) is a disk partitioning method introduced with Windows Server 2008. GPT provides: • Support for up to 128 partitions (Windows restriction). • Support for partitions larger than 2 TB. • Partitioning in Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012. With GPT, partition information is stored on the disk in well-documented locations. Each partition is identified by a GUID number. GPT also provides redundancy for the partition table. You can convert between MBR and GPT as long as there are no volumes on the disk.

A volume identifies an area of disk space that the operating system can use for storing data. A volume is either basic or dynamic, depending on the type of disk on which the volume resides.

Disk Type	Description

Basic	A *basic* disk supports volumes made up of contiguous disk space. Basic disks have limitations depending on the partition style used: • On an MBR disk, a basic disk has a limit of four partitions. You can have up to four primary partitions or up to three primary partitions and one extended partition. The extended partition can be divided into multiple logical drives. • On a GPT disk, Windows imposes a limit of 128 partitions. GPT disks do not have an extended partition. On basic disks, you can shrink or extend volumes only on the same disk. Basic volumes correspond to the primary or extended partitions on the disk.
Dynamic	A *dynamic* disk supports volumes from contiguous and non-contiguous disk space. • Dynamic disks support up to 128 volumes on both MBR and GPT disks. • Dynamic volumes allow you to extend existing volumes to include non-contiguous space on the same disk, to extend a volume onto a different disk (creating a spanned volume), or to implement striping or mirroring. • Each dynamic disk stores a replica of the dynamic disk database for all dynamic disks connected to the system. • The dynamic disk database facilitates volume management. • On a new server, import dynamic disks as a unit. The disks will be marked as *foreign*. The main reason to choose dynamic over basic disks is to take advantage of advanced disk configurations that involve spanned, striped, and mirrored volumes. It is best practice to convert a disk to dynamic only when the volumes require more than one disk.

2.7.3 Volumes

Volumes
0:00-0:02
In this video, we're going to talk about volumes.
Basic Disks
0:03-0:14
After you initialize a disk, it's going to be listed as a Basic Disk, which is the Default Storage. It's the ideal storage. We really want to stay at Basic if we can.
Partitions
0:15-2:16
Basic Disks are divided up into Partitions, and you can have up to four primary partitions. Primary partitions can be used to boot the computer. Originally, we could have up to a quadruple boot system, or you could have three primary partitions and one extended, which then would get divided up into logical drives.
One thing that you're going to see in the demos that's a little bit strange is Microsoft has dropped the word 'partition' from the operating system. Even though the Basic Disks are still divided up into partitions--primary or extended--when you actually create one, you're going to see New Simple Volume. If the disk is basic, you're actually creating a partition.
Sometimes we divide up the space on the disk because we want to isolate data. For example, on my laptop what I will commonly do is divide up Disk 0 into a C: drive and a D: drive. I can store my data on the D: drive if I need to reimage C:, because the operating system is messed up, I can go ahead and reinstall my operating system. No problem. My data is on D:. As soon as the new operating system is there, the data is there. I don't have to worry about getting it off, getting it back on.
In a server, there are many different reasons why you might want to divide up a big hard drive. The key to know is this: if it's a basic disk, once you go beyond your three primary, the computer--if you're in the GUI--is going to assume that you want an extended partition. We can't work with extended partitions directly. We have to create these logical disks. Realistically, if you need more than three drive letters out of that physical disk, you're going to be working with an extended partition, and you're going to be getting logical drives.
This is very, very rare. Usually, we do things much differently, but at least you know what the system is and how it works. If we're not working with basic disks, we might be working with dynamic disks.
Dynamic Disks
2:17-3:14
Basic disks can be converted into dynamic without loss of data. With that being said, they should only be converted to dynamic if I need to extend a volume over more than one disk.
The natural environment is a basic disk. As soon as you go to dynamic, life gets a lot more difficult. Don't go to dynamic unless you need any of the features of dynamic disks, and then we're going to talk about why you would, and you'll see this is not something you're going to be doing often in your career.
One of my first complications is this: if I take a dynamic disk and I bring it to a new server, first of all, because I'm using dynamic to tie two disks together, I should install all of the disks and import them as a unit. Each one of those disks is going to be marked 'foreign'. What I'm then going to have to do in Disk Management is Import Foreign Disks. All of the stuff we're talking about now, for the most part, happens in Disk Management. On those dynamic disks, there are different types of volumes.
Simple Volume
3:15-3:19
A Simple volume is just that. It's just a chunk of disk space.
Spanned Volume
3:20-3:37
The next type of volume we have is a Spanned volume. The first diagram I'm going to show you doesn't necessarily have to do with spanned, but I just want to kind of get this information out there when I was going to diagram some disks. It's not something that you really have to get involved with, but it's something you should be aware of. Otherwise, you could potentially run into trouble.
System Partition
3:38-4:00
When I just install Windows Server 2012, or 2008, or 2008 R2, or Vista, or Windows 7, Windows 8, anything 2008, Vista or above, it actually creates a little partition here that's invisible. You don't see it in the operating system. That's the system partition, which is where Windows stores the files necessary to boot the computer.

Boot Partition

4:01-4:27

Then if I said it could use the whole hard drive, it'll grab all the rest of the space, and that becomes the boot partition where I have the C:\windows directory.

Now, here's the important thing to remember. That partition is used to boot the computer. Whatever you do, don't go into Disk Management and get rid of it. You will be sad. You can build computers without it, but it's put in a separate partition to protect the operating system. You really shouldn't mess with that system.

Spanned Volume Continued

4:28-5:47

Now we'll take a look at our spanned volume. In a spanned volume, I had a volume, and here you can see I've got Disk 0. It's a 100 GB. I'm using 40 GB for disk C:. What I've done is I've made the rest of it, let's say, into disk F:, and right now that's worth 60 GB. We just made the numbers small to keep the math easy. Now, for whatever reason, the data in F: drive has grown beyond or needs to grow beyond that 60 GB, but your boss comes in and says, "No, we don't have time to move that 60 GB down here to another disk, because it's 2 o'clock on Friday, and everything has to get done. So you've got to fix this in the next five minutes or less or you're fired." You say, okay. How can I get more space into drive F:?

I don't have anywhere I can expand. If there was unallocated space here, I could extend the volume, and you should know that you can extend volumes. In this case, I can't do that. What I'm going to do is go down to Disk 1 and grab some space here. I grab another 40 GB on Disk 1. I create a spanned volume. This is also drive F:. What's going to happen physically is the computer is going to fill up all of the space on Disk 0, and when it gets beyond 60 GB, it's going to kind of spill into Disk 1.

Possible Disk Failure

5:48-6:12

With spanned partitions, I'm tying together space on at least two different disks. They don't have to be equal sizes, but here's the problem with this. I'm now doubly exposed to disk failure, because if either Disk 0 or Disk 1 goes down, drive F: will be lost. Both of these disks would need to be converted to dynamic before I could do this. If you just go ahead and do it, it's going to say, "Let me convert it to dynamic for you." You say, okay, and you're good to go.

Striped Volume

6:13-7:55

Our next type of volume that we have supported on dynamic disk is a striped volume. Let's take a look at what that looks like. In a striped volume, I need at least two disks, could be 2 through 32, I at least need two. I'm going to go ahead and create my volume using space from both of these disks. It has to be an equal amount of space. If I'm using 60 GB up here, then I need to be using 60 GB down here, and they're both going to get the same drive letter. What happens with the striped volume is it saves an equal amount of data to each drive. Though a striped volume with two disks, every single file, half of it will be here, half of it will be there. Striped volumes are like spanned volumes in that, as soon as I create one, I'm doubly exposed to disk failure. This is not advised in any kind of environment where you're worried about the integrity of your data.

Why would we do this? Well, the inside of a hard drive kind of looks to me like an old time record player. I've got these disks called platters, and then there's a head that reads information from the platters. If I stripe, essentially what I'm doing is, I have both of those heads reading and writing data at the same time. My reads and my writes are going to be faster. The only time we use striping is in a situation where we're really worried about speed. We're not necessarily very worried about integrity. We don't care if the disk fails. Some other mechanism in the background is taking care of fault tolerance and disaster recovery. We're just looking for as much speed as possible.

It's a performance benefit, but absolutely no fault tolerance. Our next type of volume that we can make on a dynamic disk is a mirrored volume.

Mirrored

7:56-9:09

You're going to see it's exactly what it sounds like. Sometimes this is called RAID 1. Sometimes striped is called RAID 0. Any one of those would be correct.

In a mirrored drive, it's going to be equal amounts of data. They're going to have the same drive letter. You actually won't see the drive letter for one of them. What happens with this, it's going to be exactly two disks. It writes 100% of the data on each disk. I have a perfect mirror on each disk. This is the only type of fault tolerance we can provide for the operating system. That C: drive can't participate in any of the other types of volumes except mirrored, where I could come in and I could make a mirror of my C: drive, and then I have fault tolerance.

Watch out, too, if you're taking tests that are based on multiple choice questions. Notice when I mirror, I'm mirroring by volume. Disk 0 has two volumes on the disk. If I'm being asked to mirror the entire Disk 0, I need to create two mirrored volumes, one that mirrors the C: volume, and another that mirrors the F: volume.

My last type of volume that I can create on dynamic disks we call RAID 5.

RAID 5 (Striping with Parity)

9:10-9:37

That's what's in the operating system. These are the words in the operating system. Here you may see it called striping with parity, which would also be correct. Now in striping with parity, I need at least three disks. The amount of space I use has to be identical on all three of those disks. It could be up to 32, but I certainly need at least three, and they're all going to get the same drive letter.

Parity

9:38-12:16

Now, parity is information that we can use to rebuild missing data. We always lose one drive to parity. In a striped set, I get all the space, and in a mirrored, I just get the space of one drive. Here, I just lose one drive to parity. It gets more and more efficient as you add drives. Here, I'm losing 33% percent of my disk space. If I had four, I'd lose 25%. If I had five, I'd lose 20%, and so on and so forth.

As the computer stores files to this volume--let's say we have our first file-- maybe it puts half here, half here, and then on this one, it puts the parity information. Now I go and I save my second file. I'm going to put some lines to make it easy. Maybe in this case, it puts half there, half there, and here's where it puts the parity information. I save my third file. Let's say it puts half here, half here and up here we have the parity.

You can see that even though I lose one disk to parity, it's not just one disk that's designated for parity. The parity is actually spread across the disks. Now, the great thing about RAID 5 is I can have 10 disks tied together. I can lose up to one disk and that volume will keep working. It will be very, very slow, but it will keep going. Here's what happens. Let's just say I lose disk 1. Disk 1 dies. Well, for my first file I've got half of the file, and I can use the parity to rebuild the other half.

For my second file, I've got the whole file. It'd be pretty easy to recalculate the parity. In my last file, I've got half the file here and I've got the parity, so I can rebuild the missing half. If you're looking at storage pools, striping is the same as simple, mirrored is the same as mirrored, and RAID 5 they just call parity. Speaking of storage pools, storage pools are a better option. Again, everything that we're talking about we would be doing inside of Disk Management--at least for this particular video. This is really software RAID. It's been around for quite a while, but it's not as efficient as hardware RAID. Even if you're using software RAID, it's not going to be as efficient as storage pools, which provide the same functionality, except for spanned, but give me more flexibility about adding and removing disks. You should definitely be aware of all the volume types, and make sure you're aware that they're only supported by dynamic disks, the difference between basic and dynamic and pros and cons, number of disks for each type.

In real life, it's highly unlikely you're going to do any of this.

Basic to Dynamic/ Dynamic to Basic

12:17-12:25

I can go from basic to dynamic-- no problem, no loss of data. If I want to go from dynamic to basic, I have to delete everything, and then it will revert to basic. The last thing we're going to talk about are mount points.

Mount Points

12:26-14:26

We'll look at the facts of them and then we'll see kind of what they're used for. They must be mounted to an empty folder on an NTFS volume. That's important. What we do is we use that folder instead of a drive letter to access the volume, and effectively what it does is it adds space to a specific path in the file system. Let me show you what I mean.

Here I have my C: drive and it's 40 GB. Let's say I'm going to install an application, and the application has got to be installed into the C:\App folder. It's got a hard requirement. I can't put it on D:, E:, F:, G:. No, it's got to be in C:\App. This particular application, just to make it fun, requires 60 GB--conveniently, the size that I have left over here. Now, it is possible, starting with Windows Server 2008, to both extend a volume and to shrink it.

In this case, I could probably go through and I could just extend into the unallocated space. Let's say there isn't any unallocated space here. As a matter of fact, I've got a D: drive using all 60 GB. I've got to get those 60 GB into C:\App. Here's what I'm going to do. I'm going to create an empty folder named C:\App. I'm going to create a new partition down here. This can be done with basic disks-- it doesn't have to be dynamic--worth 60 GB. Then I'm going to mount that into the C:\App folder. When I look under drive C:, let's say I've got 4 MB of space left. Let's say on the C: drive, all you've got left is 4 GB of free space. But I'll go in, I'll look in Computer, I'll click on C:, I'm going to see 4 GB free space. When I open up the C: drive and I click on C:\App, over on the right hand side, it's going to show that that folder has 60 GB of free space, but only in that one folder. That's essentially what we use a mount point for.

Those are some facts about volumes.

Summary

14:27-14:57

Make sure you know the difference between a basic and a dynamic disk. Make sure you know the five volume types for the dynamic disk: simple, spanned, striped, mirrored, and RAID 5. Make sure you know the number of disks required for each, and you should know about mount points as well, and have some idea about the file systems. That being said, you're better off staying as far away from dynamic disks as possible. What you really want to do if you're going to use software RAID is go with the storage pool.

2.7.4 Volume Formats

Volume Formats
0:00-0:07
Let's talk about the different choices we have for formatting volumes in Windows Server 2012.
Term Table
0:08-1:19
I put this table together because we're going to talk about file systems. You want to be a little bit familiar with some of the terms. I just went all the way back to the bit, which is just either 1-0 or 1-1. If I get 8 bits, then it's a byte. Whoever was doing this must have been hungry, because between, I have 2 bits, which is a nibble, and a 4, which is crumb, or vice versa. These things always make me hungry.
Because it's binary, if I get 1024 bytes all in one spot, that's my kilobytes. 1024 kilobytes, that's my megabyte. 1024 megabytes, that's my gigabyte. 1024 gigabytes, that's my terabyte. Now a days, most people are familiar up to that point in the list. We're going to see some bigger names. If I actually get 1024 terabytes, that would be a petabyte. 1024 petabytes, that's my exabyte.
Now, new at Server 2012, we're getting up even higher than that, so 1024 exabytes would be a zettabyte, and 1024 zettabytes would be a yottabyte. I think I saw recently something like this--600 exabytes of information out there in the universe. We're really looking ahead to accommodate bigger and bigger storage, but we're not quite there yet.
Volume Formats
1:20-1:25
We've got three volume formats that we can take advantage of inside of Disk Management.
exFAT
1:26-3:02
exFAT is new. I believe that came in with 2008 and Vista, although it wasn't really easy to get to in those operating systems. It's got a file size limit of 16 exabytes. It's only recommended for flash memory. If you do use it on your flash memory, it does not support ReadyBoost.
Traditionally, the very first file system was FAT, which you sometimes see referred to as FAT16. It was very limited. It didn't store data very efficiently, and so you would lose a lot of space. It just wasn't great.
Microsoft then invented FAT32, which allowed for bigger volumes and stored information more efficiently, but really, it was limited to volumes up to 32 GB, and a file size of 4 GB. They worked around that 32 GB partition limitation in some of the versions of Server. Again, they kept FAT32 around to accommodate for smaller volumes. Because NTFS has some overhead, maybe I don't need that on a 2 GB flash drive.
Now that the flash drives and the external hard drives are getting bigger, and bigger, and bigger, they've invented exFAT to take advantage of that, and use that space more efficiently, but still provide compatibility with other operating systems. This is just the latest version of FAT, has some backwards compatibility, but it supports much bigger files.
If you're a MAC person, if you format your external hard drive or your flash drive with exFAT, it's my understanding that you can use it both on PC and MAC.
NTFS
3:03-3:18
NTFS has been the file system standard for quite a while. They had one version in NT 4.0. They actually call that NTFS 4.0. With Windows 2000, it's NTFS as we know it today.
There's four main selling points.
Security
3:19-3:28
Security allows me to put NTFS permissions on the files themselves. We have a whole other video in this series that talks about NTFS security.
Encryption
3:29-3:46
Encryption lets me scramble the file, so that only the user that encrypted it or people that they designate can open it up. With encryption there is an awful lot of support that you have to provide behind the scenes, so I would say don't encourage users to encrypt files unless you'd rather have the data be lost than compromised.
Compression
3:47-3:51
NTFS compression basically allows you to save space on the hard drive.

Disk Quotas
3:52-4:04
Disk quotas lets me go ahead and limit the amount of space that the user is allowed to use on that hard drive. NTFS disk quotas are done by volume, and we track the users' usage based on the ownership of the file.
Compression Cont.
4:05-5:07
You want to be careful with NTFS compression. When I think of NTFS compression, I think of those bags they sell where you sort of vacuum pack your down blankets. That's going to save you a lot of space. When you go to get the blanket that's vacuum packed in the closet, you're going to un-vacuum pack it from that bag and take it out; it's going to take longer to grab that down blanket than the one that's just sitting in the cedar chest at the end of your bed.
I can compress files to save space in the hard drive, but it puts a burden on the processor, because the processor's got to unpack or decompress those files in order for you to use them. You also might want to be aware that certain files compress better than others. Vacuum packing a down blanket--lots of space saved. Vacuum packing a brick, not so much.
Files that are already dense like JPEGs, certain videos, pictures--you don't gain much when you compress them. Certainly on a server, if you can avoid NTFS compression, you're much better off.
ReFS
5:08-5:17
The last type of format we have is ReFS, which is new with Windows Server 2012. There's some great things about this, and some not so great things.
Pros of ReFS
5:18-6:20
Pros of ReFS, better resiliency of extremely large volumes. We're already looking at exabytes with NTFS, exFAT, so we're talking really large--larger than is available for sale right now.
The really cool thing about ReFS is, it can auto-detect data corruption. You try to run Check Disk on this and you can't even do it. It's always looking for data corruption, and it can even section off pieces of the hard drive and say, "Okay, I'm not going to store anything there, because it's corrupt." You can automatically repair errors without taking the disk off line.
In NTFS, if I want to do a Check Disk, I actually have to dismount that volume. If it's the C: volume, I have to reboot. It's not required with ReFS. Talking about how large the disks are, it supports volume sizes up to 1 yottabyte, and we saw the chart before. Why don't I do everything ReFS?
First off, it came in with Server 2012. Only Server 2012 going forward is going to be able to look at these volumes. Again, if you're concerned about Data Recovery Services, don't go ReFS.
Limitations of ReFS
6:21-6:50
Here's some other limitations. It does not support, among other things, File Compression, Disk Quotas, Encryption, Short Filenames, the System or Boot Partition.
My understanding is that eventually Microsoft's goal is to have the capabilities of ReFS match NTFS, but they're not there yet. Unless you have a partition that can't support NTFS and you've got some limitations, probably stay away from ReFS for now. It's to support the stuff that's coming, but it's not really in play.
Summary
6:51-7:05
Those are the different choices that we have for formatting volumes in Windows Server 2012: exFAT, which is great for flash drives, NTFS, which has been around for quite some time, and then ReFS, which isn't quite here, but coming up pretty fast.

2.7.5 Volume Facts

Use dynamic volumes on dynamic disks to take advantage of advanced disk configurations that provide for increased performance or fault tolerance. The following table describes the dynamic volume types:

Volume Type	Description
Simple	A *simple* volume contains a single, contiguous block of space from a single hard disk. • You can extend a simple volume onto the same disk, as long as the disk space is contiguous. • Simple volumes do not provide any performance or fault tolerance advantages.
Spanned	A *spanned* volume combines areas from two or more disks into one storage unit. The primary purpose of a spanned volume is to add more storage space to an existing volume. A spanned volume: • Fills the first area, then the second, and so on. • Does not provide fault tolerance. If one hard disk fails, you lose all data. • Cannot contain system or boot files. • Can include space from between 2 and 32 physical disks. The amount of space used on each disk can be different. • Has no overhead; all disk space is available for storing data.
Striped (RAID 0)	A *striped* volume breaks data into units and stores the units across a series of disks. When you save a single file on a striped volume, pieces of the file will be found in all disks in the array. Striped volumes: • Use two or more disks. The amount of space used on each disk must be the same. • Provide an increase in performance. Multiple disk controllers read and write data at the same time, reducing the overall time to read or write any file. • Do not provide fault tolerance. A failure of one disk in the set means all data is lost. • Have no overhead--all disk space is available for storing data.
Striped with parity (RAID 5)	A *striped volume with parity* combines disk striping across multiple disks with parity for data redundancy. Parity information is stored on each disk. If a disk fails, its data can be recovered using the parity information stored on the remaining disks. RAID 5 volumes:

	• Require a minimum of three disks, with equal space being used on each disk. • Provide fault tolerance. Data is available even if one disk in the set fails, though performance is significantly impaired. • Provide an increase in performance (although the performance is not as good as that of a striped volume). • Have an overhead of one disk in the set for parity information. A set with 3 disks has 33% overhead. A set with 4 disks has 25% overhead. A set with 5 disks has 20% overhead.
Mirrored (RAID 1)	A *mirrored* volume stores two copies of each file, with one copy on each disk or disk set. Mirrored volumes: • Require two disks. • Mirror only volumes. • Require that a mirror of a volume be the same size as the volume. • Provide the only type of fault tolerance for the operating system. • Provide fault tolerance. If one disk fails, data is preserved on the other disk, and the system switches immediately from the failed disk to the functioning disk to maintain availability. • Do not increase performance. • Have a 50% overhead. Data is written twice, meaning that half of the disk space is used to store the second copy of the data. Disk duplexing is a type of mirroring. Disk duplexing uses two hard drives and two separate disk controllers. Disk duplexing eliminates the single point of failure when a single disk controller is used.

2.7.6 Managing Disks and Volumes

Managing Disks and Volumes
0:00-0:03
In this video, we're going to take a look at managing disks and volumes.
Disk Management
0:04-0:22
To do that, we need to get into Disk Management. We can see there are some hard drives that have been added to the computer, but currently, they're offline. I'm going to bring them online. Once the disk comes online, the next thing I need to do is initialize it.
Initialize
0:23-0:44
Not all disks come in offline. The reasons these are is because I've hot added them to the machine, and that caused them to be offline. When you first add a disk to the computer, it has to be initialized. I would right click and Initialize Disk.
MBR and GBT
0:45-1:46
When you initialize a disk, you have to choose between MBR and GPT. MBR disks are limited to 2 TB. GPT disks are used if you have disks that are greater than or equal to 2 TB.
Notice it says "GPT partition is not recognized by all previous versions of Windows". This came in with Windows Server 2008 and Vista, and those are the operating systems. Any of those two or better on your client or server side should recognize GPT. If you don't have a 2 TB volume, I recommend you go MBR. Reason being, if you had to send that disk out for data recovery, there are a lot of data recovery programs that don't operate in the newer versions of the operating system, and they may not be able to help you with data recovery. As a general rule, if I don't GPT, I'm going to stick with MBR. If I need to convert it after the fact, it's certainly possible to do that.
I would just right click "Convert to GPT". As long as there are no volumes, I can go back and forth all day. Once I start creating volumes, then I'm done.
Basic Disks
1:47-1:57
These disks are all basic disks. If you're just going to work with them as individual hard drives, you should leave them basic. I can go in on my basic disk and I can make different types of volumes.
Simple Volume
1:58-2:54
A Simple Volume is just a chunk of disk space. We'll make a simple volume there's 300 MB. Assign a drive letter. Now, I have a simple volume drive F:, which is 300 MB. Notice if I need to, I can extend my volume or I can shrink it. Shrinking came in with Windows Server 2008 and Vista -- it's got to be NTFS partition and there's got to be free space, but if I need to, I can shrink it.
If I'm going to extend it, the unallocated space should be right next to this volume. We can see that there is plenty of unallocated space right next door. If there were another volume in between drive F: and the unallocated space, officially according to Microsoft what I should do, is backup that volume, delete it, expand drive F:, and then bring the other volume in from backup. A Spanned Volume uses disk space from two different disks.
Spanned Volume
2:55-4:14
What it does is fill up all of the space on one disk and then spill into the next one. We'll say on Disk 2, we're going to use 200 MB of space, but we'll also add in some space from Disk 3, and let's give it 400 MB on Disk 3.
My entire volume will be 600 MB, but it's going to be split between Disk 2 and Disk 3. What the computer will do as I add data is fill up all of Disk 2. When it gets beyond 200 MB, it would flow into Disk 3. Anytime we have two or more disks working together, we can't use basic disks anymore. We need dynamic disks. There's my spanned volume. I highly recommend against this, because now it doesn't matter whether it's Disk 2 or Disk 3 that fails. Either way the G: volume is going to go down. Spanning gives me absolutely no fault tolerance. If I wanted to convert my disk to dynamic manually, I can just do Convert to Dynamic Disk and then I'm good to go.
Striped Volume
4:15-4:55
A Striped Volume spreads data across however many disks are involved in it. You need a minimum of two. Here, it's going to be an equal amount of space, because the computer will save an equal amount of data to each of those disks. With two disks, it will put half the data on Disk 2, half on Disk 3. This is a performance benefit, but it doesn't provide me any fault tolerance whatsoever, because now, if either one of these disks fail, my striped volume will fail. If I'm looking for fault tolerance inside of Windows, I might want to create a Mirrored Volume or RAID 5 Volume.

Mirrored Volume
4:56-5:24
In a Mirrored Volume, it puts 100% of the data on both disks. Here, I'm choosing 300 from each disk, but my volume is only 300 MB, because 100% of the data will be saved to both disks simultaneously. If either one of those disks should fail, I'll be just fine, because I have a perfect mirror of it on the other disk. The last type of volume I can create in here is a RAID-5 Volume.

RAID-5 Volume
5:25-6:10
RAID-5 Volume requires at least three disks. It's also known as Striping with Parity. If I choose 300 from each disk, one of the disks is going for the parity information. You will lose one out of the disk for parity, but it spreads them across of all of the disks. It's not like one disk is going to be dedicated to parity, the other two will be data.
Parity will be spread around, but you end up losing one of the disks. So with three disks, I lose 33% of my space, with four is 25%, with five, 20%, so the more disks I have in a RAID-5 array, the more efficiently I use the space. Here, I can lose up to one disk and the drive will keep functioning. It's going to be very slow, but it will keep functioning.

Storage Pool
6:11-6:33
All of these types of volumes, while they're supported under Microsoft ... there's really much better ways of doing this. I wouldn't necessarily recommend any of these -- not spanned, striped, mirrored, RAID-5. It's much better if you're looking to use software RAID to use a storage pool and manage it that way rather than trying to make up fault tolerant volumes.

Mount Point
6:34-8:08
The last thing I'm going to show you is using a mount point. Let's use drive F as an example. Let's say in drive F: we're going to install an application. Drive F: has 300 MB. Let's say our application requires 500. I can't expand drive F: because there's no unallocated space next door, but let's say we know this application is going to install itself into the F: App folder. I only need to provide the 500 MB into that particular folder. I don't need to expand the entire F: drive. For that, I would use a mount point.
I'm going to go ahead and create a Simple volume 500 MB, but I'm not going to assign it to drive letter or path. You can mount it right from here. Every once in a while, that throws an error. So I'm going to do it as a separate step. Now, I have a new volume down here, it's worth 500 MB. It doesn't have any drive letter. I've no way of accessing it right now, but I'm going to right click it, Change Drive Letter and Paths, and I'm going to add my mount point, which is over on F:, and it's the App folder. It must be on an NTFS partition. The folder must be empty. If I go in now and I look at App, you can see the folder now has a little shortcut icon and I've got 500 MB just in that particular folder. If you need to add space into one folder, that's done with a mount point.

File Systems
8:09-9:04
The only other thing I will show you are my file systems. I actually have access to a number of file systems. FAT and FAT32 are for backwards compatibility. You would probably not be using those. NTFS has been the gold standard for Microsoft for quite a while. ReFS is new with Windows Server 2012. ReFS has support for very large volumes. It doesn't have all the functionality of NTFS quite yet. They're looking to expand it as time goes on. Unless you're going to create volumes and have disks that go beyond the limits of NTFS, right now you don't necessarily want to go with ReFS. These are huge volumes that they support. The hardware isn't really out there quite yet. We're going to be expanding that direction as time goes on, but not quite now. That's how we work with volumes in Disk Management.

2.8 Storage Pools

As you study this section, answer the following questions:

- What are the advantages of the VHDX format compared to the VHD format?
- What are the advantages of a dynamically expanding disk compared to a fixed disk?
- Which virtual hard disk type provides the best performance?
- What are the advantages of storage pools compared to traditional storage?
- What is thin provisioning?
- What is contained in the primordial storage pool?
- What are the disk requirements for adding a disk to the storage pool?

After finishing this section, you should be able to complete the following tasks:

- Create and mount a VHD.
- Create storage pools.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 1.0 Configure Windows Servers.
  - Configure Server Storage
    - Create and Mount Virtual Hard Disks (VHDs)
    - Create a Storage Pool

This section covers the following 70-410 exam objective:

- 103 Configure local storage.
  - This objective may include but is not limited to:
    - Design storage spaces
    - Create and mount virtual hard disks (VHDs)
    - Configure storage pools and disk pools
    - Create storage pools by using disk enclosures

2.8.1 Creating and Mounting a VHD

Creating and Mounting a VHD
0:00-0:14
In this video, we're going to take a look at creating and mounting virtual hard disks, and we can do this right through Disk Management. I'm going to access this through the Tools menu; Tools, Computer Management, and we'll go ahead and click on Disk Management.

Create New VHD
0:15-0:22
To create a new VHD, we'll just go right up the Action menu, create VHD. We need to give the VHD a name.

Give the VHD a Name
0:23-0:31
If you don't type the .vhd the wizard will actually just go ahead and add that.

Pick a Size for the VHD
0:32-1:07
My next step is to pick the size. I'm going to make mine pretty small. Notice it's using the VHD format, new with Windows Server 2012. We could specify the VHDX format if I needed a VHD larger than 2040 GB in size, and that will support up to a maximum of 64 TB. It's also a little bit more resilient for power failures. The only problem is it wouldn't be supported in anything earlier than Windows Server 2012. So if you have any idea that you might need to bring this back to an earlier operating system and you don't need that size, you might as well go VHD.

Fixed Size or Dynamically Expanding
1:08-1:59
The last choice I'm going to make is between fixed size, which they recommend, and dynamically expanding. Dynamically expanding means that the file for the VHD is only going to get as big as the things that I put inside of it. Fixed size, on the other hand, immediately is going to grab all of that space in the hard drive. Why would Fixed size be recommended? With Dynamically expanding, as I add items to the VHD file, the processor has to keep making that file bigger, so best performance would be Fixed size. It's also good because you know, right away, that that amount of space has been used up in the hard drive, and you're not in a situation where the VHD is going to grow, and use up all the space on the physical hard drive and crash the server. Because all we're doing is looking at an example, I'm going to choose Dynamically expanding, but again, you see that Fixed size is recommended. Once I've created my VHD, it comes in exactly as if I've attached a new hard drive to the computer; so I have to go ahead and Initialize it, just like I would a regular disk.

Initialize
2:00-2:12
Then, once it's been Initialized, I have to partition it and format it again, exactly the way I would with a brand new hard drive.

Partition and Format
2:13-2:40
The computer has alerted me that I have to format it, but that's actually been done automatically. If I look inside a computer, I'm going to see my new drive, and here's drive E:.

Detaching and Reattaching VHD
2:41-3:11
Now with a VHD, if I reboot the server, it won't be attached anymore, so we're going to unattach our VHD and reattach it just so you can see what that process would look like. I want to right click my VHD and Detach VHD. Now it's not attached. The file itself still exists, anything I've put inside it still exists, I just don't have access to it anymore through the operating system. If I had rebooted the server, I'd be in exactly the same situation.

Attach an Existing VHD
3:12-3:57
To attach an existing VHD, just go right back up to the Action Menu, Attach VHD. I can either browse or type the path to my VHD. Notice I have the option to attach it as Read-only; that means I'd be able to look at whatever's inside the VHD, but I wouldn't be able to make any changes. If I need to be able to make changes or delete things inside of it, I should not attach it as Read-only. If for some reason I needed to Change the drive Letter, I can do that just as I would with a regular physical disk. That's how we create, attach, and manage virtual hard drives using Disk Management.

2.8.2 VHD Facts

A virtual hard disk (vhd) is a file that is created within the host operating system that simulates a hard disk for the virtual machine.

- The virtual machine accesses the virtual hard disk through the management operating system.
- When you create the virtual disk, you specify the disk and controller type (such as ATA/IDE or SCSI).
- The virtual disk type does not have to match the physical disk type. For example, you can create a virtual SCSI disk on a physical ATA disk.
- Do not create a virtual hard disk on a folder with encryption enabled. You can create a virtual hard disk on a volume with BitLocker enabled.

You have the following options for configuring a hard disk for use by the virtual machine:

Option	Description
Virtual hard disk format	The vhd format is the traditional format. The vhd format: • Supports files up to 2,040 GB in size. • Provides compatibility with previous versions of Windows. The *vhdx* format is a new format introduced in Windows Server 2012. The vhdx format: • Supports virtual disks up to 64 TB in size. • Is more resilient during power failures. • Offers enhanced compatibility with large sector disks. • Is supported only in Windows Server 2012.
Virtual hard disk type	A *fixed* disk occupies a set amount of hard disk space in the management operating system. The size of the virtual hard disk file indicates the total storage capacity of the virtual disk. • This disk type takes longer to create than other disk types. • The entire disk size, including any empty space within the virtual hard disk, is reserved on the physical disk. • Performance is improved because the entire virtual disk is a contiguous block and because additional space does not need to be allocated to expand the size of the disk. For best performance, use fixed disks for all virtual disks. A *dynamically expanding* disk allocates physical disk space in the vhd file as virtual disk storage increases. • Physical disk space is allocated only as it is used by the virtual machine.

	• The size of the vhd file grows as more disk space is used by the virtual machine. • This disk type makes the most efficient use of hard disk space. • It is possible for the physical disk to run out of space as the vhd file size grows.

Keep in mind the following when creating a virtual hard disk:

- After creating the virtual hard disk, you must initialize, partition, and format it.
- The vhd detaches when you reboot the server.
- You must manually reattach the vhd whenever you reboot or restart the server.
- To manage a vhd file from the command line, use **diskpart**.
- When creating a vhd file using **diskpart**, enter the vhd size in megabytes.

2.8.4 Creating Storage Pools

Creating Storage Pools
0:00-0:54
In this video, we're going to look at creating storage pools. To create a storage pool, we need to go into File and Storage Services. In here, I should see any hard drives that I have installed in my computer, and then I can go over and make a Storage Pool.
Any drives that I have not assigned to a storage pool that I created, show up in the Primordial storage pool until I assign them to a pool that I create. Essentially, what storage pools let me do is work with these hard drives as if they were using hardware RAID. It's the best way to describe it.
When I combine the disks into a storage pool, they're going to appear as if they were one disk to the operating system when in fact they're multiple disks being managed. The reason I do this is to give myself some flexibility with the hard drives so that I can add hard drives to the pool and take them away without actually bringing down the disks and volumes that are being stored on that pool.

Storage Pool in Disk Management
0:55-1:11
Before we create our pool, I just want to take a look at what it looks like in Disk Management. We're going to go into Computer Management. You can see, right now, from the operating systems perspective, all four of these disks are available. It's going to change once we create our pool.

Create the Pool
1:12-1:25
To create the pool, we're going to up under Tasks and click New Storage Pool. We give the storage pool a name, and then we select which disks will participate in it.

Allocation: Automatic or Hot Spare
1:26-2:28
Each disk we specify can have either an Automatic allocation or be used as a Hot Spare. If it's Automatic, the space from that disk will be available into the Storage Pool. If it's a Hot Spare, that space would not be used for disks and volumes. It would be kept in case one of the other disks fails. We'll do one Hot Spare and two Automatic, and then create our pool.
I've got a storage pool with three members: two Automatics, one Hot Spare. Disk 4 remains in the Primordial storage pool. If I look in here, all I see is Disk 4. I don't see the members of my storage pool, because I haven't created a disk on that storage pool yet. That would be my next step. We'll start the New Virtual Disk Wizard.

Layout
2:29-2:32
When I create my disk, I have three layouts.

Simple
2:33-2:41
A simple layout is the equivalent of striping. The data will be spread across the disks, but if one disk fails, there'll be no fault tolerance.

Mirror
2:42-2:56
Mirror is the equivalent of a mirrored volume. It will store data on at least two disks. If one fails, I'll be able to keep going. Because I have a hot spare, if one failed, it would immediately bring the hot spare in and then rebuild that mirror.

Parity
2:57-3:15
Parity requires at least three disks. What it does is stripe data and parity across the disks such that if one disk fails, I'll still be able to keep going. Make sure you know what these layouts do and how many disks they require.
We're just going to do a simple disk.

Provisioning Type
3:16-3:18
Now I'm going to choose whether it's going to be thin or fixed.

Fixed
3:19-3:34
Fixed is also sometimes called thick. If it's fixed, it means it's going to use space from the storage pool up to the volume size. It's going to grab that space from the pool right away, and then my disk can only be as big as the space that I have in the pool.

Thin

3:35-5:21

Let's say, for example, I've got my pool now, and I know the pool's got 400 GB in it, but I've got a few more disks on order. Within the next month or so, I actually should have 700 GB worth of disks in my server. I can actually use thin to specify this particular disk as being larger than the amount of space I have now available. You've got to make sure that that volume doesn't grow beyond the actual physical space that you have available. If you do need to do that for some reason, it's absolutely fine.

You can specify size. If I'm doing thin, the assumption is, I am specifying a size greater than the amount of free space I currently have so I need to type that in. If we do fixed, in that case, I can come in and just say "give me the maximum size", because I'm limited to the amount of space that's in the storage pool.

It also tells you up here that if you choose fixed and you're using either mirror or parity, it's going to take up a little bit more space than is actually there. You want to go through and either manually leave some space. Easiest thing to do is just say, "give me the maximum size".

Now that I've added a disk -- I'm going to cancel the New Volume Wizard -- it's going to appear to the operating system. If I go into Disk Management, you can see Disk 5 is actually my storage pool. All my disks are 127 GB. Here's Disk 4, which is not part of any pool.

Now, because I have a pool with two disks and one hot spare, my pool is 251 GB. Two of those disks that are set to automatic. The hot spare is not included, because it's just meant to be a spare if one of those other two disks fails.

Creating a Volume on the Storage Pool

5:22-6:10

Now I would go through and create a volume on the storage pool exactly the same way that I would with any disk. I can either do it in Disk Management, or I can do it out here by right clicking and making a New Volume.

If you need to get rid of the virtual disk, I would detach it before you delete it. Notice, I also have the ability to extend it so that later on, if I buy more hard drives, then I can just add them into the pool, which is a great thing about storage pools, whereas if we were doing some type of RAID in computer management, we wouldn't be able to extend that. Once you create that volume, RAID 5 or mirrored volume, essentially, you're done. You could also come in here and create my volume in here as well, just the way I would any volume.

Summary

6:11-6:41

Storage pools allow me to manage my disks as one unit, so that the entire pool appears as one volume to the computer. I can create multiple virtual disks on it, so that it would look as if they were multiple virtual disks. I'm only going to see the virtual disk that I create in Disk Management. I'm not going to see the members of the pool, because they're being managed as a pool.

This gives me flexibility with my storage and a lot better fault tolerance than the software RAIDs that were available before Windows Server 2012.

2.8.5 Disk Enclosures with Storage Pools

Disk Enclosures with Storage Pools
0:00-0:17
In this lesson we're going to look at several enhancements Microsoft has added to storage spaces in Windows Server 2012 R2. Specifically we are going to look at storage tiers and we're going to look at enclosure awareness. Let's begin this lesson by talking about storage tiers.

Storage Tiers
0:18-0:41
Storage tiers are designed to help system administrators who have to manage file servers. The key problem that file server administrators have is the fact that the file server is expected to store a lot of stuff and it's expected to be very fast, therefore file server administrators want the best performance they can out of their file servers. They want their file servers to run wicked fast.

Solid State Drives (SSD)
0:42-1:27
Now one way to meet this need is to replace the SCSI or SATA drives in a server with solid state drives. Now solid state drives use flash storage technology instead of the typical hard disk platters. This makes them much, much faster than traditional hard disk storage technology.
However if you've ever shopped for an SSD drive then you know they have two key weaknesses. Number one, they are rather small by way of comparison to traditional hard disk drives. They just don't store as much. And they also cost a lot more than a comparably sized typical hard disk drive. So they're fast but they're smaller and they're a lot more expensive.

Stale Data
1:28-2:03
The key problem here is the fact that end users tend to utilize file servers as a catch all location to just dump stuff--whether they ever intend to use that stuff again or not. The result is that we end up with a file server that has a ton of data being stored on it that may or may not even be accessed. As a system administrator you don't know which of that data is good and which of that data is stale and will never ever be used again. Therefore it actually isn't all that cost effective to invest in SSD drives to store all of this information.

Classify and Migrate Data
2:04-2:46
What we want to do is classify our information somehow based upon how frequently it's used. What we would like to do then is migrate any information that is being used a lot to our fast SSD drives and any old information that hasn't been touched in three or four years we want to migrate to our cheaper storage devices such as a SATA drive or maybe a SCSI drive.
Now classifying this information manually would be--well not even possible. It would be so labor intensive as to not be feasible. However, with Windows Server 2012 R2 we can use storage tiers in order to accomplish this automatically.

Storage Tiers
2:47-2:58
Now essentially what storage tiers do is allow you to create a single virtual disk that is actually comprised of two different tiers of storage devices.

SSD Tier
2:59-3:20
For example we may have an SSD tier that's used only for data that's accessed frequently. This is the data that needs to be very fast. For example, maybe it's year-end and our accounting department needs quick access to all of the files they need to create their year-end reports. That data needs to be over here on an SSD drive.

Hard Disk Drive Tier
3:21-4:04
Likewise we can create a second tier within the same virtual disk that uses a standard hard disk device, maybe a SATA drive, for example. This drive contains stuff that maybe people don't use very often.
For example, maybe there is a design document used by the Research and Development department that they used to create a product say three years ago and we moved on since and we're developing a new version of that product. Therefore that old design document, while still important and does still need to be accessed from time to time, we're not looking at it every day, therefore we don't need to put it on the SSD drive. Instead we can just put it over here on our slower and less expensive SATA hard disk drive in the server.

Migration is Transparent to the User

4:05-5:06

The cool thing about storage tiers is the fact that data is moved transparently by storage spaces between the different tiers based upon how frequently the data is accessed. Basically storage spaces is going to be looking at data that's been used a lot and automatically migrated over here to the SSD drive, and any data that's not being used very often is going to be automatically migrated down here to our SATA drive.

The key thing to remember is that all of this is completely transparent to the end user. The end user sees a virtual disk. They have no idea that it's being composed of multiple storage devices with differing performance characteristics. All they see is a single virtual disk over here.

So why is this important? Well, essentially what storage tiers do is combine the best attributes of SSD drives and the best attributes of standard hard disk drives--all combined into a single storage space and a single virtual disk.

Tracks How Frequently Data Chunks are Accessed

5:07-5:22

In order to accomplish this, what tiering does is track how frequently the data on our virtual disk is accessed. It dynamically promotes and demotes data based upon what it's called its I/O temperature.

Hot Data vs. Cold Data

5:23-6:18

Hot data is data that is used more frequently and therefore it's going to be automatically migrated to our SSD drive. Likewise any data that is not used very frequently is classified as cold data and it's going to be migrated automatically behind the scenes to our less expensive traditional hard disk drive. By doing this our frequently used data is on the fastest most expensive storage medium, while our infrequently accessed data is stored on our slower less expensive storage media within the same virtual disk.

Now I will point out that there may be situations where you have a piece of data that is important and does need to be stored on the SSD tier, but maybe isn't accessed frequently enough to be classified as hot data, in which case it would be automatically stored over here on our slower performance, less expensive storage here.

Manually Pin Data in the SSD Tier

6:19-6:41

In that situation you can actually manually pin that data and say no matter how frequently it's accessed, it still needs to be stored in the SSD tier because I know at some point it's going to be needed and it's going to be needed in a bad way and we need to have the best performance possible. Even though it doesn't meet the criteria, migrate it to the SSD tier anyway.

Scheduling a Tiering Task

6:42-7:22

Now the tiering feature is a scheduled task that's going to run every day at 1 AM by default. During the day no migrations take place. We don't want to impact performance of the server. We're going to wait until one in the morning when probably nobody is using the system and then we'll run our migration based upon classifications at that point.

Now you can modify the schedule manually if you need to. If you run a 24/7 shop you may have to pick a different time than 1 AM in order to run this, so you can decide when that best time would be. The key is to make sure that you migrate the data only when system utilization is at its lowest point.

Storage Tier Requirements

7:23-8:05

With that in mind let's talk about what you need to do in order to create a storage base using storage tiers. The first thing we need to look at are the requirements for doing this. First of all, the storage pool must have a sufficient number of hard disks and SSD drives to support whatever storage layout that you want to use. The disk must have enough free space in order to do this.

In addition, if we're going to create a storage space with storage tiers the virtual disk must use fixed provisioning. That's because volumes that you create on the virtual disk that you storage tiers have to be the same size as the virtual disk itself.

Building a Storage Tier

8:06-9:36

With this in mind let's take a look at an example of building a typical storage tier implementation.

Now in this example we're anticipating that about one third of our data is going to be classified as hot and roughly two thirds of our data is going to be classified as cold. Therefore when building the storage tiers I'm going to use say four SSD drives here, but we're going to use eight standard hard disk drives, probably SATA drives in this storage pool. Again, about one third of the data we anticipate being hot and about two thirds of the data that's going to be stored in this pool that we're going to assume it's going to be cold.

In order to build the storage tier we will take this available storage space on all 12 of these devices and allocate them to a single storage pool. Then within that pool we're going to create two different tiers. In this example we labeled them the SSD tier, and the HDD tier. Again, the hot data is going to be automatically migrated to the SSD tier which is composed of the SSD drives, and the cold data is going to be automatically migrated to the HDD tier which is composed of our standard SATA disk drives.
These two tiers will then define our single virtual disk. And once defined our users can begin dropping files into the various shares on the server and the migration process will run at one in the morning, by default, and move data to the appropriate tier.

Storage Tier Cmdlets

9:37-9:45
Now there are several new PowerShell commandlets available for working with storage pools that use storage tiers. You need to be aware of these.

Set-FileStorageTie

9:46-10:06
The first one is called Set-FileStorageTier. This is used to assign a file to a particular tier. In essence, we are pinning the file to a particular tier. We're overwriting the classification process and saying, "No, this file always needs to in the HDD tier or the SSD tier--no matter what the classification rules may say."

Get-FileStorageTier

10:07-10:16
Next we have the Get-FileStorageTier commandlet that this gets the files that are currently assigned to a storage tier on the volume and gets their status as well.

Clear-FileStorageTier

10:17-10:22
Next we have the Clear-FileStorageTier. This removes a file from a particular tier.

New-StorageTier

10:23-10:26
We have the New-StorageTier commandlet which creates a new storage tier.

Get-StorageTier

10:27-10:33
We have the Get-StorageTier commandlet which gets the various storage tiers that are defined on a storage space.

Set-StorageTier

10:34-10:38
We have the Set-StorageTier commandlet which modifies an existing storage tier.

Resize StorageTier

10:39-10:45
We have the Resize StorageTier commandlet which increases the size of a storage tier.

Remove-StorageTier

10:46-10:50
We have the Remove-StorageTier commandlet which simply removes a storage tier from the storage pool.

Get-StorageTierSupportedSize

10:51-11:00
Then finally we have the Get-StorageTierSupportedSize commandlet which gets the minimum and maximum possible sizes of the storage tier.

Enclosure Awareness

11:01-11:31
Now the last thing we're going to look at in this lesson is the concept of enclosure awareness. Now in Windows Server 2012 storage spaces include an enclosure awareness feature. This enclosure awareness feature will mirror data between multiple enclosures. The idea being that if one enclosure fails or goes offline, then the data will remain available in one or more of the alternate enclosures.

Requirements
11:32-12:28
Now in order to use enclosure awareness with storage spaces you have to meet several key requirements. First of all, your storage enclosures must support SCSI Enclosure Services or SES. If they don't, then you can't use enclosure awareness. Then you need to implement the appropriate hardware and the key point--as noted here--is the fact that different tolerance levels require different amounts of hardware.
For example, if you need to tolerate one failed enclosure with two-way mirrors then you actually need three separate compatible storage enclosures. By way of comparison, if you need to tolerate two failed enclosures with three-way mirrors then you actually need five compatible storage enclosures. In essence, think of enclosure awareness as RAID 1 mirroring for entire disk enclosures instead of individual disks.
Summary
12:29-12:48
That's it for this lesson. In this lesson we introduced you to several of the new features and storage pools in Windows Server 2012 R2. We first looked at storage tiers where we can classify data based upon how often it's used and move it to faster or slower storage accordingly. Then we looked at storage enclosure awareness where we mirror physical disk enclosures.

2.8.6 Storage Pool Facts

Storage spaces are logical drives for storing data and other user files. Storage spaces are created from *storage pools*, a collection of space from multiple disk drives or other storage devices. A storage space displays and is managed as one drive regardless of the number of disks or devices contributing space to the storage pool.

Storage spaces are composed of three components:

- *Devices* are the hard disks or other types of storage from which storage pools are created. You can use a variety of devices to create storage pools, such as SATA drives, SCSI drives, and external USB drives.
- *Pools* of storage are created from the available disk space. A pool is composed of the free space available on the specified storage devices.
- *Storage spaces* define the virtual disks created from a pool. One or more storage spaces can be created from the pool. To the Windows system and the user, storage spaces appear as disks with typical drive letters (E drive, F drive, etc.).

Storage spaces can be managed dynamically and they eliminate the need for such tasks as repartitioning drives, resizing volumes, and backing up data in order to repartition. When you need more disk space for your storage spaces, follow these steps:

- Install a new storage device to the system.
- Add unallocated space on the device to a storage pool.
- Allocate space to an existing storage space.

A disk must be online and unallocated (not formatted or partitioned) to be added to a storage pool.

The following table identifies configuration options in storage pool creation:

Option	Description
Allocation	Each disk in the pool is allocated as automatic or as a *hot spare*. • Space allocated as automatic becomes available in the storage pool. • Space allocated as a hot spare is reserved for use in the event that a disk in the pool fails.
Storage layout	Storage spaces can include data resiliency implemented through storage layout options. When you create a virtual disk, choose a storage layout option. The options include: • *Simple*, which does not provide redundancy. This option simply adds space from the storage pool to the storage space. When you select the **Simple** option, all of the data in the storage space is lost if one of the drives fails. This option is similar to RAID 0 (data striping). • *Two-way mirror* requires at least two storage devices.

	The data is written to two devices. Two-way mirror requires twice as much device space as the amount of storage allocated to the storage space. This option protects you from a single storage device failure. This option is similar to RAID 1 (mirroring). • *Three-way mirror* requires at least five storage devices. The data is written to three storage devices. This option provides redundancy for the data if two storage devices fail at one time. • *Parity* requires that you have at least three storage devices. This option uses parity information to reconstruct data if one of the storage devices fails. Parity uses less space for redundancy than the mirror options, but performance is not as good as the mirror options if a device failure occurs. Parity requires only 50 percent more redundancy space than storage space. This option is similar to RAID 5 (parity striping). All storage layouts except simple use more free space than their specified size.
Provisioning	*Fixed provisioning* allows you to specify the size of the space based on actual disk size. You can use all of the space on a disk or part of the space on a disk. *Thin provisioning* or *overbooking* allows you to allocate larger storage spaces than disk space available in the pool. • Thin provisioning is based on the premise that not all users will use all of the space in their allocated storage space. • Space is added to a user's storage space as the user consumes space. • If a storage space runs out of disk space, it will immediately unmount, leaving any I/O processes vulnerable to data corruption. An unmounted storage space must be brought back online manually. Files can be accessed after the storage space is brought back online manually, but you must add more physical disk space to the pool and add it to the storage space in order to use the storage space. Windows creates the virtual disk only if there is sufficient free space.
Storage tiers	In Windows Server 2012 R2, storage tiers allow you to combine solid-state drives (SSDs) with hard disk drives (HDDs) in the virtual hard drive of storage spaces. Based on file usage, storage spaces automatically determine the type of drive on which the data is stored. When using storage tiers, storage spaces: • Store frequently accessed data on an SSD for quick access. • Store less frequently accessed data on an HDD. • Analyze usage daily or based on a configurable timeframe. • Evaluate data in 1 MB sized pieces for tier assignment.

	• Allow administrators to assign (pin) a file to a specific tier. Storage tiers require fixed provisioning. Storage tiers also require the number of columns to be the same on both tiers.
Write-back cache	In Windows Server 2012 R2, storage spaces can leverage solid-state drives within a storage pool to create a write-back cache. This cache performs two key functions: • It protects against data loss in the event of a power failure. • It buffers small, random writes to solid-state drives that are very fast, but have less capacity. The cached data is then written to slower, high-capacity traditional hard disk drives at a later time when disk utilization is low. Latency within the storage pool can be significantly reduced when using a write-back cache. A 1 GB write-back cache is enabled by default when a new storage space is created if the storage pool contains enough physical SSD disks. To use write-back caching, the storage pool must meet the following requirements: • A simple storage space must have at least one SSD drive. • A two-way mirror storage space must have at least two SSD drives. • A three-way mirror storage space must have at least three SSD drives. If there aren't enough SSD disks in the pool, then the write-back cache size is automatically set to **0**. This effectively disables write-back caching on the pool. Write-back caching is compatible with storage tiers. However, be aware that when the write-back cache is enabled in a storage space that uses storage tiers, the fast storage tier will actually be 1 GB smaller than the size you configure because of the SSD space allocated to the write-back cache.

Keep in mind the following about storage spaces:

- The *primordial* storage pool contains unallocated disks that are connected to the server but are not assigned to a storage pool.
- Once a storage space is created, you create one or more volumes on the drive and manage the disk as you would a physical disk.
- Use the **Detach Virtual Disk** option to detach the storage space before deleting it.
- Automatic rebuild in Windows Server 2012 R2 automatically rebuilds storages spaces from storage pool free space in the event of drive failure. Automatic rebuild eliminates the need for hot spares in storage pools as long as there is a sufficient number of drives assigned to the pool.

Limitations of storage pools include:

- Disks containing the boot volume, system partition, or any Cluster Shared Volumes (CSV) cannot be added to the pool.
- The individual drives in the pool must be at least 10 GB.
- A storage pool must have at least one drive.

- Storage pools using Fibre Channel and iSCSI are not supported for failover clusters.
- Storage pool virtual disks to be used with a failover cluster must use NTFS.

PowerShell commands you can use to manage storage spaces include:

- **New-StoragePool** creates a new storage pool.
- **Add-PhysicalDisk** adds a new disk to an existing storage pool.
- **New-VirtualDisk** creates a virtual disk in the storage pool.
- **Get-StoragePool** retrieves information about the storage pool.

Enclosure awareness, introduced in Windows Server 2012 R2, provides an added level of fault tolerance in which each copy of data is associated with a particular JBOD (just a bunch of disks) enclosure. If one of the disks fails, another copy of the data exists on one or more enclosures. Be aware of the following regarding disk enclosures:

- The **-EnclosureAwareDefault** parameter of the **New-StoragePool** cmdlet identifies the default allocation policy for virtual disks created in an enclosure-aware storage pool.
- Enclosure awareness mirrors data between two or more enclosures.
- Using multiple SAS adapters per enclosure and per server allows multipathing.
- Cluster Shared Volumes I/O redirection is supported on enclosure-aware disks. This mitigates data loss when the connection to the disk enclosures is lost, but the server maintains connection to the network.
- JBOD storage enclosures must support SCSI Enclosure Services (SES).
- The following table identifies storage requirements:

Failed enclosures	Configuration	Storage required
1	Two-way mirroring	3 compatible storage enclosures
2	Three-way mirroring	5 compatible storage enclosures
1	Dual parity	4 compatible storage enclosures

- The following table identifies the JBOD count and failure coverage for the specified configurations (all configurations are enclosure-aware):

Configuration	Two JBOD	Three JBOD	Four JBOD
Two-way mirroring	1 disk	1 enclosure	1 enclosure
Three-way mirroring	2 disks	1 enclosure + 1 disk	1 enclosure + 1 disk

Dual parity	2 disks	2 disks	1 enclosure + 1 disk

3.1 Virtual Machines

As you study this section, answer the following questions:

- What are the hardware requirements for a Windows Server 2012 running Hyper-V?
- How does Hyper-V in Windows Server 2012 differ from Hyper-V in previous Windows Server versions?
- What is the purpose of the Memory weight setting? How is it used?
- What is an advantage to using dynamically expanding memory?
- What are some advantages to using fixed memory?
- In what situation is smart paging used?
- What is the difference between resource metering and resource control?

After finishing this section, you should be able to complete the following tasks:

- Create a virtual machine.
- Manage a virtual machine.
  - Add and manage hardware
  - Change boot order
  - Change memory allocation
  - Manage processors
  - Configure Integration Services
  - Configure shutdown options

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 2.0 Hyper-V.
  - Manage Virtual Machines with Hyper-V Manager
    - Create and Manage Virtual Machines

This section covers the following 70-410 exam objective:

- 301. Create and configure virtual machine settings.
  This objective may include but is not limited to:
  - Configure dynamic memory
  - Configure smart paging
  - Configure Resource Metering
  - Configure guest integration services
  - Create and configure generation 1 and 2 virtual machines
  - Configure and use enhanced session mode

3.1.1 Virtual Machines

Virtual Machines
0:00-1:03
In this video, we're going to talk about Virtual Machines.
The role that you need to install for virtual machines is Hyper-V, and once you've got that installed, you can create virtual machines. Really the idea behind this was, we have a server, it's got some great hardware, but maybe we're not using all that hardware to the maximum potential. Yet, on the other hand, I don't want to put one copy of Server in there, then load it up with all these different roles, because maybe they possibly conflict with each other, or it's easier to troubleshoot if they were running on different servers. I can create virtual machines, which allows me to run multiple operating systems on the same physical hardware.
New with Windows Server 2012, we have some features we can set up when we're creating or configuring the virtual machines. When you go through and create the virtual machines, you're going to specify how much RAM is available to that virtual machine, and that will be taken from the physical RAM for the Server. You need to make sure that you allocate some for the host, the computer that's running Hyper-V, otherwise you're going to run into a problem. We can set up the processors a little bit, and we've got some stuff we can do with the network cards. There's a lot of new functionality with 2012.

Dynamic Memory Requirements
1:04-1:37
The first thing we're going to talk about is Dynamic Memory.
It requires that your host be running Windows Server 2008 R2 Service Pack 1., or Server 2012. The guest operating systems also have to be able to work with Dynamic Memory. That would be guests running 2003 SP2, 2008 SP2, 2008 R2 SP1, and then even some of the clients will work with this. Windows Vista, it's got to be Ultimate or Enterprise, and it has to be SP2. With Windows 7 -- again Ultimate or Enterprise -- it has to be SP1. Then we also have support for Windows 8.

Dynamic Memory
1:38-2:24
The concept behind Dynamic Memory is this: the server has a certain amount of RAM available to it. Let's say my server has 32 GB of RAM. As I create virtual machines, I can assign them a certain amount of startup RAM. In the past, before Dynamic Memory, once I assign memory to that virtual machine, that's it. That's how much RAM it has, it's taken away from the host, it's not available to other virtual machines, and if I want to change it, I had to shut the virtual machine down. With Dynamic Memory, I can specify a range of memory that could be allocated to that machine. I can also go through and tweak it so Hyper-V knows which machines should get extra memory if more than one machine needs extra memory. Of course, all of the machines put together, it can't go beyond the physical memory of the host, also subtracting whatever the host needs to run.

Dynamic Memory Settings
2:25-2:26
Let's look at some of our Dynamic Memory settings.

Startup RAM
2:27-2:33
When you initially create your virtual memory through the wizard, you're going to assign startup RAM. That's the memory required to start the virtual machine.

Minimum RAM
2:34-2:54
Now, virtual machines need a little bit more RAM during startup than they actually need for performance. You can set a Minimum RAM, which is the minimum amount of memory to be assigned to the VM, and that could even be less than the startup memory. If your startup memory is 1024 MB, but down here it's 512 MB because it works fine with 512 MB once it's booted, Maximum RAM is the maximum amount of memory the VM is allowed to use.

Maximum RAM
2:55-3:02
Hyper-V won't allocate beyond that amount.

Memory Buffer
3:03-3:40
We can also set a memory buffer, which is how much memory is assigned to the VM compared to the amount of memory actually needed by the applications and services running inside the VM. It basically holds a certain percentage back. It looks at how much memory the virtual machine actually needs and divides that by the buffer value over a hundred. I'm not great with math, but let's say you say, "Put a 20% buffer, well divide that over a

hundred, and then it will go into this, and basically it will grab 20% of how much RAM the machine actually needs to hold as a buffer. Make sure that if it needs more RAM it's already there. If there's not enough physical memory for all the machines, it doesn't bother with the buffer, the buffer is not maintained.

Memory Weight

3:41-4:06

I can also set a memory weight, which determines how to distribute the memory among the VMs if not enough physical memory is available in the host to give every VM it's requested amount of memory. Presumably, don't set them all to the same weight. Important VMs, set it to a higher weight -- less important machines, a lower weight. They're still all going to get their minimum amount of RAM, but this is when they can't get everything that they need. If this Hyper-V has to make any decisions, you're telling it how to weigh that decision.

Smart Paging

4:07-5:06

There's also a feature new with 2012 called Smart Paging. Basically the idea is this: If I need a certain amount of memory to boot my VM, but it doesn't require as much to run it, and maybe I've dropped down to the minimum memory, if I have to restart that VM, there might not be enough physical memory to bump it back up to the startup memory. What the computer does is this; it's going to provide a reliable restart if they're running with less than the minimum startup memory, by making a little temporary paging file on the hard drive. On the hard drive it makes a temporary memory, a little paging file, uses that to reboot the virtual machine, and then it gets rid of that little paging file. It's only used when a virtual machine is being restarted, and there's no available physical memory to get it back up to that startup memory. No memory can be reclaimed from other virtual machines running on the host. It's not going to rob other machines on the host just to reboot one of them. On the other hand, if you don't have the startup memory when you're booting the VM for the first time, it's not going to let you boot. The startup memory has to be there when you first boot it. Smart Paging is just for reboots.

Resource Metering

5:07-6:25

Also new with Server 2012 is something called Resource Metering. Basically, what this does is allow the administrators to track statistics for billing. It's enabled on a per virtual machine basis, and here's your PowerShell command. Enable-VMResourceMetering, pretty much the obvious command. You can also monitor and generate reports on virtual machine usage, which would be used to bill clients. You could use the Measure-VM PowerShell commandment, the -VMName switch, and give the name of the virtual machine, or I could Get-VM, name of the virtual machine, Measure-VM, and then select all. It will show me all the resources that are being metered, and what the values are. Basically, the idea behind resource metering is, your company is selling out the use of Hyper-V. It's really coming into Support Cloud Computing, and Cloud Computing is where you have some servers, or you may or may not have some servers at your location, but when you need more processing power or more server power than you actually have at your company, you can hop up to the cloud, and you sort of pay as you go. If you were hosting cloud computing, you need to be able to determine how much resource your client used on that VM so you can properly bill them, because again, it's pay as you go. This is all done in PowerShell. We're not going to see any of this in the GUI when we take a look at Hyper-V.

Resource Control

6:26-6:32

Resource metering is different than Resource Control. We would see resource control in the properties of our processor.

Percentage of Total System Resources

6:33-6:44

We've got percentage of total system resources, which is the percentage of VM time. It's measured by how many processors are assigned to the physical computer, so we're setting up how much of the total resources of the host this virtual machine can use up.

Relative Weight

6:45-6:57

Relative Weight tells Hyper-V how to decide how CPU is distributed. So, a virtual machine with a higher weight, let's say 500, is going to get twice the CPU time as a virtual machine with a lower weight, let's say, 400.

Integration Services

6:58-7:05

We also want to go ahead and talk about Integration Services. This essentially allows the virtual machine to interact with the host.

Drivers that Enable the Guest OS to Interact with the Host

7:06-7:39

The first thing that it does is, provides drivers that enable the guest operating system to interact with the host, and to use the host's hardware. If you're having hardware issues inside the virtual machine, you should reinstall the Integration Services. So, you got a flaky network driver, the video driver's misbehaving, try reinstalling the Integrations Services. From the Action menu of the VM, you're going to Insert Integration Services Setup Disk. It just puts an ISO in the virtual DVD drive, and then you're going to install it like regular software. Now, that's Integration Services inside the VM.

We also have some settings in the Properties of the virtual machine itself.

Operating System Shutdown

7:40-7:56

The first one we're going to see is Operating System Shutdown. This basically allows the host machine to gracefully shut down the virtual machine if the host is being rebooted. So, it can send the shutdown command to the guest operating system, the VM shuts down, and then the host shuts down, so that there're no crashes with the VM.

Time synchronization

7:57-8:05

Time synchronization allows the VM to take its time from the host. You want to be very, very careful with time. I'll talk to you about that in a second.

Data Exchange

8:06-8:23

Data exchange is responsible for exchanging management information between the host and the guest. This will be things like the Host Name, Operating System Version, Build Number, and other things. All of these are checked by default, and you can leave them checked, but you should know what they are and that they're in that integration services area.

Heartbeat

8:24-8:46

We also have support for a Heartbeat. Basically a Heartbeat means the computer sends out a packet -- it's usually like every second -- and what happens is if the host doesn't get that Heartbeat, it knows that the virtual machine might have locked up or crashed. Then it can initiate failover or live migration, whatever you have set up, to account for the fact that the Hyper-V guest machine has crashed. But this is how it knows whether the guest is running or not.

Backup (volume snapshot)

8:47-9:16

You also have an option in there that says Backup (volume snapshot.) We know that files that are open can't be backed up by backup software. So, Volume Shadow Copy supports that. It lets it make a snapshot of the file, and it'll back up the snapshot. You have to have this turned on to enable online backups of the Hyper-V virtual machine. Otherwise, the virtual machine would need to be taken off line for a backup of the host to be performed. Again, it's checked by default, but any of this is fair game for any exam you might want to approach in your future.

Time Synchronization

9:17-9:56

Time synchronization. I just want to say something about that. Time is really critical in a network. If times are off, then you're going to have weird errors. I say weird because I can't be specific. It will be sporadic, and sometimes they're crazy. Usually the messages don't have anything to do with what the problem is. So, make sure the time zone in your virtual machine is synchronized with the host, and if you're going to leave time synchronization on, make sure that they get synchronized with the year. Anytime you have something really weird going on, stop for a minute, check time, because if time isn't right, you're going to have all kinds of problems.

Those are some of the new features that we have for creating and managing virtual machines in Hyper-V.

3.1.2 Virtual Machines on 2012 R2

Virtual Machines on 2012 R2
0:00-0:17
In this lesson, we're going to spend some time looking at several of the enhancements that have been added to Hyper-V virtual machines in Windows Server 2012 R2. Specifically, we're going to look at VM generations. We're going to look at Quality of Service, and then we're going to look at Enhanced Session Mode policy.

Virtual Machine Generations
0:18-1:05
Let's begin by looking at virtual machine generations. The generation that's assigned to a virtual machine will determine what kind of virtual hardware will be assigned to the virtual machine and also the functionality that will be provided to the virtual machine. Windows Server 2012 R2 introduces Generation 1 and Generation 2 virtual machines. In previous versions of Windows Server, there was no generation distinction for virtual machines. In fact, all virtual machines running on Windows Server 2012 and earlier are just considered Generation 1 virtual machines. In fact, sometimes you'll hear these older virtual machines referred to as simply traditional virtual machines, so if you hear the term traditional virtual machine, we're talking about Generation 1 virtual machines.

Generation 1
1:06-1:18
On Windows Server 2012 R2, you can still use Generation 1 virtual machines if you want to. If you do, the virtual machines will use the same virtual hardware that was used in previous versions of Hyper-V.

Generation 2
1:19-1:49
Generation 2 virtual machines, on the other hand, are only available on Windows Server 2012 R2. Generation 2 virtual machines use a simplified virtual hardware model and they support the Unified Extensible Firmware Interface, or UEFI, instead of the traditional BIOS-based firmware. In addition, a lot of the legacy virtual hardware devices that we used in traditional virtual machines have been removed from Generation 2 virtual machines.

Requirements
1:50-2:42
Let's take a look at Generation 2 virtual machines now in a little more detail. Let's begin by looking at the requirements for using Generation 2 virtual machines. First of all, when creating Generation 2 virtual machines, you have to use Windows Server 2012 R2 as your Hyper-V hypervisor. In addition, you're somewhat limited as to which guest operating systems you can install in your Generation 2 virtual machines. You can create Windows Server 2012 virtual machines, you can create Server 2012 R2 virtual machines, you can create a 64-bit Windows 8 virtual machine, or a 64-bit Windows 8.1 virtual machine. If you want to create a Server 2008 virtual machine or, say, a Windows 7 virtual machine, then you've got to use a Generation 1, or traditional, virtual machine instead of Generation 2.

Generation Assignments Cannot be Changed
2:43-3:23
There's one key point you've got to keep in mind when you define your virtual machine and set its generation, and that is the fact that once you set that generation, you cannot change it. If I create, say, a Server 2012 virtual machine and I configure it to be a Generation 2 virtual machine and I install my operating system in it and later just, oh, I want to go back and make this a Generation 1 virtual machine because I want to use a piece of legacy virtual hardware that's not available in Generation 2, you're out of luck. You can't do it. If you want to use legacy hardware in your virtual machine, then you've got to set it up as a Gen 1 virtual machine instead of Gen 2 at the outset. You can't go back and forth between them.

Generation 2 Features
3:24-3:29
With that in mind, let's take a look at some of the new features in Generation 2 virtual machines.

PXE Boot
3:30-3:58
First of all, you can PXE boot now from a Generation 2 virtual machine using a standard network adapter or a synthetic network adapter. In previous versions of Hyper-V, if you wanted to perform a remote installation of the guest operating system via PXE Boot, you were required to install a legacy network adapter, and these legacy network adapters are no longer available in Generation 2 virtual machines. Generation 2 virtual machines support PXE Boot using our standard network adapter.

SCSI Boot
3:59-4:26
Generation 2 virtual machines also support SCSI controller booting. In previous versions of Hyper-V, you couldn't boot a virtual machine from a SCSI-attached virtual hard disk or DVD. Generation 2 virtual machines, on the other hand, can boot from a virtual hard disk or DVD that's attached to the SCSI controller. In fact, you'll see that the virtual IDE controller is no longer available in Generation 2 virtual machines.
Secure Boot
4:27-4:48
Generation 2 virtual machines also support Secure Boot. Secure Boot is a feature that helps prevent unauthorized firmware or operating systems or UEFI drivers from running at boot time. Secure Boot is actually enabled by default for Generation 2 virtual machines. However, you can modify this once the virtual machine has been created.
UEFI Firmware Support
4:49-5:27
As we've alluded to several times, Generation 2 virtual machines provide UEFI firmware support, which replaces our traditional BIOS. The question comes up at this point, if I want to use UEFI firmware in my virtual machine, does my physical hardware have to also use UEFI? The answer to that is no. UEFI firmware is actually not required on the physical host in order to use UEFI in a virtual machine. The virtual machine firmware and its configuration are completely independent of the physical hardware that the virtual machine is running on.
Generation 2 Limitations
5:28-5:32
With that in mind, let's take a look at some of the limitations of using Generation 2 virtual machines.
RemoteFX
5:33-5:52
First of all, Generation 2 virtual machines do not support RemoteFX. RemoteFX is a set of technologies produced by Microsoft that are designed to make the visual experience associated with Hyper-V VMs a lot nicer. You can't use RemoteFX in Gen 2 virtual machines.
Virtual Floppy Disks
5:53-5:59
Much to your sadness, I'm sure, you can no longer use virtual floppy disks with Generation 2 virtual machines.
Legacy Network Adapters
6:00-6:09
And as we talked about just a second ago, the legacy network adapter is no longer available. Only synthetic adapters (standard adapters) are used.
Quality of Service Management
6:10-6:40
With that in mind, let's take a look at Quality of Service management. This is another new feature of Generation 2 virtual machines. Using storage Quality of Service, or QoS, as we call it, we can control the throughput of data to virtual disks. In Windows Server 2012 R2, you can use Hyper-V Manager to define the minimum and the maximum input/output operations per second value for each hard disk in a VM.
Enhanced Session Mode Policy
6:41-7:57
Next, let's look at Enhanced Session Mode policy. Windows Server 2012 R2 introduces Virtual Machine Connection Enhanced Session Mode. Enhanced Session Mode is a feature that allows you to redirect local resources to a virtual machine session. If you've used Hyper-V, you know that in previous versions you had to initiate a Remote Desktop Connection to a virtual machine in order to redirect local resources from the host to that virtual machine. For example, let's say you needed to copy and paste some text between the virtual machine and the host operating system. In this situation, you actually couldn't do it using Hyper-V Manager. Instead, you had to enable Remote Desktop on the virtual machine and then you had to establish Remote Desktop connection with the virtual machine, and then you could move data back and forth.
This introduced a key problem. For example, let's say I was testing a security patch and my virtual machine was running within a sandboxed environment with no network connectivity either to the host or to the external network. In that case, I was out of luck because I had to have a Remote Desktop connection in order to transfer data back and forth between the virtual machine and the host.

Local Resources Can Be Redirected

7:58-9:02

In Generation 2 virtual machines, using Enhanced Session Mode, you can now actually redirect local resources using a Desktop Connection session using the Virtual Machine Bus instead of a network connection.

Basically, what we're doing with Enhanced Session Mode is removing the requirement to have a network connection on the virtual machine in order to redirect our local resources, but be aware that Remote Desktop Services still has to be running on the virtual machine. We still rely on Remote Desktop to redirect local resources. We're just no longer requiring a network connection to use Remote Desktop. We're using the Virtual Machine Bus instead. Using Enhanced Session Mode, several different types of local resources can be redirected: we can redirect the audio, we can redirect printers, we can redirect the clipboard so we can copy and paste information back and forth. That's something I use a lot. Smart cards can be redirected to the virtual machine now from the local host. We can redirect USB devices.

Support for Enhanced Mode Connections

9:03-9:25

Enhanced Session Mode is really cool. The one key thing you've got to keep in mind, however, is the fact that only Windows Server 2012 R2 and Windows 8.1 guest virtual machines will actually support Enhanced Mode connections. If you have an earlier version of Windows Server or an earlier Windows Client, then Enhanced Session Mode connections will not be supported and won't work.

Summary

9:26-9:52

That's it for this lesson. In this lesson we introduced you to several of the new features available in Windows Server 2012 R2 Hyper-V virtual machines. We first talked about Generation 1 and Generation 2 virtual machines, and we identified the differences between them. We then looked at quality of service parameters we can apply to our virtual hard disks within a virtual machine. Then we ended this lesson by looking at Enhanced Mode connections that allow us to redirect local resources from the host to the virtual machine.

3.1.3 Creating Virtual Machines

Creating Virtual Machines
0:00-0:04
In this video, we're going to see how to create virtual machines using Hyper-V.
Opening Hyper-V Manager
0:05-0:22
We're on the host machine right now. We're going to go up and open up Hyper-V Manager, so I'm going to go Tools, Hyper-V Manager. You want to right click your Server and make a New Virtual Machine. It says, "Welcome to the Wizard," you say, "Yep, Next."
Give It a Name
0:23-0:46
First you're going to give it a name, and you can see here, I can opt to Store the virtual machine in a different location. Just make sure that you have a location that's got enough space. You can kind of see it grayed out here, and if we start taking snapshots of the virtual machine, it's going to be stored wherever that virtual machine file lives, so take that into account when you're creating the virtual machine, and make sure you have plenty of hard drive space if you're going to need to take snapshots.
Initial Startup Memory
0:47-0:51
We're going to set the initial Startup memory. We can change any of this after the fact.
Dynamic Memory
0:52-0:57
I can also tell it to turn on Dynamic Memory for this machine. I'm going to leave it 'til after I've created the virtual machine.
Virtual Switch
0:58-1:05
Here I can select a virtual switch for my virtual machine to be connected to. If I leave it not connected, I can change that after the fact.
Virtual Hard Drive
1:06-1:42
Now I'm going to select a virtual hard drive. By default, it's creating a new virtual hard disk. If I just use it through the wizard, it's going to be a Dynamically Expanding VHDX. That's 127 GB. If I want a fixed size VHD, I need to do that through a different wizard, so if you don't want a Dynamically Expanding VHD, then either pre create it and use an existing virtual hard disk, or say, "I'm going to make one later," and create it after the fact. If you do it through the wizard, it's going to be Dynamically Expanding, or you can convert it later. Lots of options with Hyper-V. We're just going to say, Next.
Install An Operating System
1:43-2:26
Now it asks me to Install an operating system, and I like to pick and Install an operating system later, but certainly you could specify a DVD right off the bat, either a physical drive, or an ISO file. I don't know of any operating systems that install from floppies, but you have the option, and then you can also install it from a network based server, assuming you've got connectivity.
I didn't select any virtual networks, so that's grayed out. If I had connected my network adapter, then that would be colored in and I could select this. It would be something like a WDS server, something like that. You're not going to be able to do it through this wizard anyway, so I'm going to install my operating system later--Next, and then Finish, and there's my new virtual machine.
Settings of the New Virtual Machine
2:27-2:47
I'm going to right click the virtual machine and go into Settings, and you can see that I can change any of the choices that I had made during the wizard. I can make a different virtual hard disk, I can specify what's in the DVD Drive, I can connect my network adapter, I can change any of those choices.

Installing Using an ISO

2:48-4:21

A lot of times we end up installing virtual machines from ISOs, and ISOs are just images of CDs or DVDs, and that's what we're going to be using. We're going to do an Image file, I'm going to Browse, and we'll do Server 2012. There's my ISO file, and I hit Open. Essentially, what I've done is the equivalent of opening up a real CD ROM drive on a physical machine and putting in that Server 2012 DVD. It's just that this is a virtual machine, and that's a file that's got a copy of the DVD. Hit OK.

This virtual machine has nothing on its hard drive, but it's got the Windows Server 2012 DVD in the CD ROM drive, so if I double click it, I can open up the machine. If you don't like double clicking, you can right click and hit Connect. This first button is to send Control Alt Delete to the machine. If you're a keyboard person, it's Control Alt End, e n d. Those are the keystrokes to send it into the virtual machine. This is the Start button, which I'm going to go ahead and click. Turn Off is the equivalent of hard crashing the machine, holding down the power button until it crashes. Shutdown sends the shutdown command to the operating system and it will turn off, and the Save button is like hibernating a laptop. It stops using the physical resources, but it maintains the state of the machine.

Now we've come through, and we would go through and install our operating system, and then we'd be ready to go with our virtual machine.

Creating a Virtual Machine in Windows Server 2012 R2

4:22-6:27

Let's take a look at creating virtual machines in Window Server 2012 R2. I'm going to right click my server and do a New Virtual Machine and you're going to see this wizard looks very similar to 2012--just minor differences. We'll go ahead and click "Next", as usual we give it a name. I'm going to call mine "TestG1".

Here's the biggest difference, I choose Generation 1 or Generation 2 and you can see right here once you create the virtual machine, you cannot change its generation. We'll do one with Generation 1, you still assign the memory, you still connect it to a virtual network, create a hard disk, and then choose whether to install an operating system. So, the Generation 1 these are the traditional virtual machines that all the other generations of Hyper-V have available.

Generation 2 is not going to be dramatically different in terms of creating it. So I'm going to right click and make a New Virtual Machine and this will be "TestG2". This time I choose Generation 2. You can see it doesn't change our wizard in any way, so I'm going to go ahead and hit Next, assign it memory, a virtual network adapter if we need one. Create a hard disk, choose whether or not to install an operating system and you can see that this looks a little bit different than it did with Generation 1.

I install have the option to install an operating system later but now I can use a bootable image file or a network based server. It doesn't look quite the same as it did in Generation 1, but it's still essentially the same options are available to us. I'm going to hit "Next" and then "Finish".

You can see creating virtual machines with 2012 R2 is really not dramatically different than Window Server 2012.

Summary

6:28-6:45

The only difference being that we can select the generation of virtual machine. Now what we select during the creation is going to change our options when it comes to actually managing the virtual machines.

3.1.4 Managing Virtual Machines

Managing Virtual Machines
0:00-0:22
In this video, we're going to take a look at managing the settings on existing virtual machines. I need to get into Hyper-V Manager. Go to Tools > Hyper-V Manager. There are going to be some settings that you can only change when the machine is off, so we're going to work with this Windows 8 virtual machine.

Windows 8 Adding Hardware
0:23-0:35
Now, first of all, if I need to add any hardware, I can do it from this screen. I need to add some network adapters, fibre channel adapter, SCSI controller. BIOS lets me set the boot order, and you can see by default, first it's going to try the CD ROM, then it's going to try the hard drive.

BIOS
0:36-0:52
Then, if it doesn't find anything on the hard drive, it'll go to the network adapter, and finally the floppy. If you needed to change the boot order like you would on a physical machine, you could do it in here.
In this page, I can change how Hyper-V deals with the memory.

How Hyper-V Deals with the Memory
0:53-1:26
Startup RAM is the amount of RAM it's going to get at startup, and unless I turn on Dynamic Memory, that's all the RAM this virtual machine is going to have. The virtual machines need a little bit more RAM at startup than they do to actually run. So startup memory, I'm basing it on the startup, but I might actually be able to set a minimum memory that's less, and the machine would be just as happy with that. Windows 8, Server 2012, those should work just fine with 512 MB of RAM.

Dynamic Memory
1:27-1:54
If I enable Dynamic Memory, what it does is, it allows the host to allocate memory based on need. If I set a minimum RAM of 512 MB, and there's another virtual machine that needs RAM, it might take away some of the RAM from this particular computer and give it to that other virtual machine, because it's actually doing a lot of work and this one is not doing much of anything. This has to be turned on before you boot the machine. You're not going to be able to adjust it and completely turn it on if the machine's already booted.

Minimum RAM
1:55-2:00
Minimum RAM is the minimum amount it can take it down to. Make sure that you have enough for the machine to run.

Maximum RAM
2:01-2:04
Maximum RAM is the maximum amount it can allocate to that machine.

Memory Buffer
2:05-2:18
Memory buffer says, basically, it's going to try to keep a certain percentage of the desired RAM--of the startup RAM--available for this particular machine. If there's not enough physical RAM on the host, it won't be able to keep that buffer.

Smart Paging
2:19-2:58
Hyper-V also supports something called Smart Paging if I have a machine that's rebooting. So if we're looking at this machine, my startup RAM is 1024 MB. That's actually got to be available when the machine first starts, but let's say it's been running for a couple of days and Hyper-V drops it down to 512 MB, because there's something else that needs that physical memory, and now I reboot this machine.
The machine needs 1024 MB in order to reboot, but it's already allocated that physical memory somewhere else. This machine only has access to the 512 MB of RAM, so it'll create a little special pagefile on the hard drive of the host, and use that to get the extra 512 MB that it needs to boot, and then when the virtual machine is done booting, it lets that pagefile go.

Weight
2:59-3:25
I also can come in and set a weight, to prioritize this particular machine. Let's say this machine has got 1024 MB of RAM, it's doing a lot of work, it would run better with 2048 MB. I say, well, yeah, but this is a Windows 8 machine. In my list of priorities, it's pretty low. I've got another server that also might need to increase its RAM, and I'm going to set that as high, so that if Hyper-V has to make a decision, it's going to give RAM to the higher priority machine before it will give it to the lower priority machine.

Processors
3:26-3:41
I can come in and I can set the number of virtual processors for this machine. I might have a quad core CPU in the host, but I only want to give the guest machine one core. I can only go above one if I have more than one core in the physical machine.

Resource Control
3:42-4:03
I also have resource control here, where I control the resources. I can say, Well, we're going to keep a certain percentage of the host processor in reserve for this. I can set a limit in terms of what it can use for the processor and a relative weight. Something with 200 weight versus something with 100 weight; the 200 weight will get twice as much attention from the processor as the 100 weight.

Hard Drive
4:04-4:08
Now I've got my hard drive. I can change that if I need to. I've got my CD ROM.

CD ROM.
4:09-4:15
I can put an ISO in, or I can connect this to the physical drive. SCSI Controllers allow me to add SCSI devices.

SCSI Controllers
4:16-4:22
By default, all I can add is a hard drive, but the great thing is, I can add it while the virtual machine is running.

Virtual Network/Virtual LAN/Bandwidth Management
4:23-4:35
I can change which virtual network this adapter is attached to. I can also come in and enable virtual LAN and bandwidth management. I've got some COM ports that I can map.

COM Ports
4:36-4:38
I could even put in a virtual floppy.

Virtual Floppy
4:39-4:46
Here's the name of my virtual machine. Integration Services sets up which services it's going to use, so I can enable the shutdown using that shutdown button, Time Synchronization, Data Exchange.

Integration Services
4:47-4:57

Heartbeat
4:58-5:09
Heartbeat is used for failover, new with Windows Server 2012. You can do failover right in Hyper-V without having to enable any clustering, and then whether or not this is going to support volume shadow copy backup.

Snapshot File
5:10-5:13
The Snapshot File identifies where my snapshots will be kept.

Smart Paging File Location
5:14-5:19
Smart Paging File Location: if I need to use that smart paging, this is going to set up where it would be kept. Automatic Start: what do I want this virtual machine to do when the physical computer starts?

Automatic Start

5:20-6:01

By default, if it was running when the physical machine stops, it'll restart when the physical machine reboots. Or I can have it always start automatically, or I can do nothing, which means I'd have to manually start it.

Then, this is what's going to happen if I just shut down the physical server. By default, it's going to save the virtual machine state, so that the virtual machine doesn't crash. It's probably your best option; turn off the virtual machine. That's a hard crash of the virtual machine; that's not a great option. Send it the shutdown command? Maybe so, but usually we just leave it on Save the Virtual Machine State, which is an awesome option.

Those are the configuration changes that you can make for existing virtual machines.

Windows Server 2012 R2

6:02-6:07

Let's take a look at managing virtual machines in Windows Server 2012 R2.

Settings for the Generation 1 Machine

6:08-6:52

First we'll take a look at our Generation 1 machine and what you're going to see if we right click. If we right click and go into Settings is that it's really no different than 2012. We have a BIOS where we can set up our boot order. We have memory, processor, hard drive is attached to an IDE controller.

If I go up to 'Add Hardware', I can add a 'RemoteFX 3D Video Adapter' if I had one in my computer. I can add a 'Legacy Network Adapter'. I've got my network adapter. I've got a diskette drive but no diskette in it and my integration services, etc.

So, nothing different there with Generation 1.

Settings for the Generation 2 Machine

6:53-8:04

It becomes very different with Generation 2; so let's go into the 'Settings' for Generation 2.

You can see right away there's a lot less settings. If we look at our 'Add Hardware', we don't have anything with the floppies, we don't have RemoteFX. Here, we don't have a BIOS anymore, we have the firmware, although I surely can go through and adjust the boot order. You'll notice that 'Enable Secure Boot' is also turned on and that's the one that helps prevent anything nasty from running at boot time.

I just have a SCSI controller because now we can boot off of a SCSI hard drive or DVD drive. If I expand my SCSI drive and I click on 'Advanced Features', I can 'Enable Quality of Service management' where I can set the minimum input or output per second for the SCSI drive.

I also can enable 'hard disk sharing' which lets me have a hard disk that's shared by several machines. We only have synthetic network adapters and you can see they're the same as for Generation 1 but that's the only type that we can use.

Those are some of the differences that we see in a Generation 2 machine.

Extended Session Mode

8:05-9:29

The other thing that we'll see that's different is the 'extended session mode'. If I go in to a virtual machine that's running Windows Server 2012 R2, I double click it and I connect to it, it simply connects to that virtual machine. In order to enable this, you've got to right click your server and go into Hyper-V Settings and there's your Enhanced Session Mode Policy.

Once you turn that on, and click 'OK', then Hyper-V is going to request an Enhanced Session Mode Policy connection to the virtual machine. If I double click it again, it comes up and it says, "Hey, we can control some of the display options." I can also come in and add in local resources. If I click 'more' I can add in support for Plug and Play, USB drives, that type of thing. When I click 'Connect' it doesn't look any different, but now I would have support for USB, smart cards, that type of thing, which is going to expand the ability of my virtual machine to use the hardware that's present on the physical host.

Those are some of the differences that we'll see with Windows Server 2012 R2 in terms of actually managing the virtual machines.

3.1.5 Virtual Machine Facts

Virtualization is using software to emulate one or more physical components in a computer system. Be familiar with the following types of virtualization:

- *Server virtualization* allows multiple instances of server operating systems on a single physical computer. With server virtualization, you can migrate servers from older hardware to newer computers or add virtual servers to computers with extra/unused hardware resources.
  The physical machine is called the *host* operating system.
  The virtual machine is called the *guest* operating system. A virtual machine is also referred to as a VM.
- *Network virtualization* allows multiple virtual servers to communicate using networking protocols as if they were attached to a physical network.
- *Storage virtualization* partitions physical storage on one system for use by multiple virtual servers. You can also use virtual storage to create or imitate storage devices such as iSCSI storage units.

The Hyper-V role in Windows Server 2012 enables you to create and manage virtual machines. Be aware of the following Hyper-V details:

- Hyper-V can be installed on a full installation of Windows Server 2012 or a Server Core installation.
- When you add the Hyper-V role using Server manager, the hypervisor software and the hypervisor management tools are installed.
- It is best practice to install only the Hyper-V role on a server. You can create one or more VMs on which you can install other server roles.
- Consider installing Hyper-V on a Server Core installation to minimize resources used for the Hyper-V partition.
- When you install Hyper-V on Server Core, you must manually install the hypervisor management tools.
- The Hyper-V Manager console allows you to create a VM, import a VM, configure switches, and perform other management tasks.

Windows Server 2012 running Hyper-V has the following hardware requirements:

- A 64-bit processor that includes hardware-assisted virtualization
- A system BIOS that supports the virtualization hardware
- Hardware-enforced Data Execution Prevention (DEP)

The following table describes new or updated features available in Windows Server 2012 Hyper-V and Windows Server 2012 R2 Hyper-V:

Features	Description
Dynamically expanding memory	Dynamically expanding memory maximizes the amount of memory a VM is allowed to use. Dynamically expanding memory allows you to: • Specify a range of memory that can be allocated to each virtual machine.

	• Indicate machines to which extra memory can be allocated. The total amount of memory allocated to the host and all virtual machines cannot exceed the amount of physical memory available. • Set memory for: **Startup RAM**, the memory required to start the virtual machine **Minimum RAM**, the minimum amount of memory to be assigned to the virtual machine **Maximum RAM**, the maximum amount of memory available to the virtual machine **Memory Buffer**, the memory assigned to the virtual machine compared to the amount of memory needed by the applications and services running in the virtual machine A buffer is created only if there is enough physical memory available. **Memory weight**, the distribution of memory among virtual machines
Smart paging	Smart paging allows a virtual machine to restart when there is not enough available memory to restart the virtual machine. • Smart paging uses disk resources as additional, temporary memory. • A Smart page is created when: The VM is being restarted. The restart may be a result of the host restarting. There is no available physical memory. No memory can be reclaimed from other VMs on the host. • The Smart page is deleted as soon as the VM has restarted.
Resource metering	Resource metering measures the assigned resources and the actual usage of the assigned resources on a virtual machine. Resource metering: • Allows an Administrator to track VM statistics for billing purposes. • Is enabled using the **Enable-VMResourceMetering** command. • Monitors and generates reports on VM resource usage. Use the **Measure-VM** cmdlet to monitor usage on a VM. For example, to collect usage statistics for the VM named **Sales1**, enter: **Measure-VM -VMName Sales1** You can retrieve a list of all the resources being measured and then pipe the list into a **Measure -VM** command to determine usage of those resources values. For example, to

	get usage values for all resources being used on the **Sales1** VM, enter: **Get-VM Sales1 | Measure-VM | select ***
Resource control	Resource control allows you to analyze each VM's usage of host resources. Resource control uses: • A percentage of total system resources, measured by how many processors are assigned to the computer. • Relative weight to determine how CPU resources are distributed between the VMs. A higher weight indicates that the VM should receive more resources.
Integration Services	Integration Services allows the VM to interact with the host system. Integration Services provides: • Drivers that enable the guest OS to interact with the host and the host hardware. If you are having hardware issues with the VM, re-install Integration Services using the Integration Services Setup Disk. • Settings to control how the VM and host machine interact: **Operating system shutdown** allows the host machine to gracefully shut down the VM when the host is being rebooted. **Time synchronization** allows the VM to set its time based on the host time. **Data Exchange** allows the host and the VM to exchange management information. **Heartbeat** sends a signal from the VM to the host that allows the host to determine if a VM is working properly or if a VM has locked up or crashed. **Backup**, referred to as a volume snapshot, enables online backups of the Hyper-V VM.
Enhanced Session Mode	On Windows Server 2012 R2, Enhanced Session Mode allows you to redirect local resources to a virtual machine session. In Enhanced Session mode, you can redirect resources using a Remote Desktop Connection session using the virtual machine bus. This eliminates the need for a network connection. Resources you can redirect to the virtual machine include: • Smart cards • Clipboard • USB devices • Audio

	• Printers To enable Enhanced Session Mode, right-click the server in Hyper-V Manager, select **Hyper-V Settings**, and then select the **Allow enhanced session mode** checkbox under Enhanced Session Mode Policy. Windows Server 2012 R2 and Windows 8.1 are the only guest operating systems that support enhanced mode connections.

Hyper-V on Windows Server 2012 R2 introduces generation 1 and generation 2 virtual machines. All legacy virtual machines running Windows Server 2012 and earlier are now referred to as generation 1 virtual machines. The following table identifies key components of generation 2 virtual machines:

Component	Description
Supported guest operating systems	The following guest operating systems are supported on generation 2 virtual machines: • Windows Server 2012 R2 • Windows Server 2012 • 64-bit versions of Windows 8.1 • 64-bit versions of Windows 8
Hardware	Hardware support in generation 2 virtual machines include: • Virtual SCSI controllers and virtual SCSI CD-ROMs (these replace IDE controllers and CD-ROMs). • UEFI firmware replaces legacy BIOS. UEFI provides Secure Boot. • Synthetic network adapters replace legacy network adapters and provide IPv4 and IPv6 network boot. • There is no floppy controller support in generation 2 machines. • The universal asynchronous receiver/transmitter (UART) controller requirement for COM ports has been replaced by optional UART support for debugging. • Software-based components replace the components listed below. i8042 keyboard controller PS/2 keyboard PS/2 mouse S3 video Software-based components use minimal resources and make the guest operating system more secure. • The following components are no longer required:

	PCI bus Programmable interrupt controller (PIC) Programmable interval timer (PIT) Super I/O device
Supported features	Generation 2 virtual machines support the following features: • PXE boot using a standard network adapter • Boot from a SCSI virtual hard disk • Boot from a SCSI virtual DVD • Secure Boot • UEFI firmware support
Non-supported features	The following features are *not* supported in generation 2 virtual machines: • RemoteFX • Attached physical CD or DVD drive • Attached virtual hard disk in VHD format • Boot from VHDX file converted from a VHD file • Setup mode for Secure Boot

Be aware of the following regarding generation 2 virtual machines:

- You use wizards and property sheets in Virtual Machine Manager (VMM) to specify the generation of a virtual machine.
- You can run generation 1 and generation 2 virtual machines together.
- For generation 2 virtual machines, boot order for devices is handled by the **FirstBootDevice** parameter with Windows PowerShell commands.

  By default, the virtual hard disk is the first device to boot. If the virtual machine has more than one virtual hard disk, the disk marked **Contains the operating system for the virtual machine** is booted first.

- If you use a .vhd file format to create the virtual machine or the virtual machine template, the virtual machine is set up as a generation 1 machine.

  You cannot change the generation of a virtual machine once the virtual machine is created.

3.2 Virtual Machine Storage

As you study this section, answer the following questions:

- Why is it necessary to compact dynamically expanding virtual disk files?
- Under what conditions should a virtual disk configuration not be edited?
- What is a key difference between using IDE and SCSI virtual disks?
- How is a pass-through disk implemented in a virtual machine?
- Under what circumstances would it be advantageous to use differencing disks?
- How does Hyper-V in Windows Server 2012 protect the integrity of domain controllers running as virtual machines?
- When should you create a snapshot of a virtual machine?
- What is the difference between snapshots and differencing disks?

After finishing this section, you should be able to complete the following tasks:

- Create a virtual hard disk.
- Install a parent virtual machine using a fixed virtual hard disk.
- Create differencing disks from a parent disk.
- Create, apply, and delete snapshots.
- Convert, compact, merge, and expand virtual hard disks.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 2.0 Hyper-V.
  Manage Virtual Machines with Hyper-V Manager
  - Create Virtual Hard Disks (VHDs)
  - Create Differencing Drives (Parent-Child)

This section covers the following 70-410 exam objective:

- 302. Create and configure virtual machine storage.
  This objective may include but is not limited to:
  Create VHDs and VHDXs
  Configure differencing drives
  Modify VHDs
  Configure pass-through disks
  Manage snapshots
  Implement a virtual Fibre Channel adapter

3.2.1 Virtual Hard Disks

Virtual Hard Disks
0:00-0:14
In this video, we're going to talk about virtual hard drives in the context of Hyper-V. When you create a virtual machine, you need a VHD virtual hard drive to store the operating system and the data on. There are some new changes and options with Windows server 2012 Hyper-V.

.VHD and .VHDX
0:15-0:40
Traditionally, the format for virtual hard disks are .VHD. They support virtual hard drives up to 2040 GB. New with server 2012, we have a .VHDX format, which is for larger drives up to 64 TB. Supposedly, these have greater resilience for power failures as well.
The only reason you might have to stick with VHD is if it needs to be backwards compatible with some type of version of Hyper-V that doesn't support VHDX.

Copy the Contents of an Existing VHD to a New VHD
0:41-0:58
Another really neat thing that you're going to see when you're creating a VHD is, there's an option to just copy the contents of an existing VHD to the new VHD. You could make a VHD that has a generic install of server 2012. Have that hanging around every time you make a new machine. Just copy the VHD, and you don't have to install server 2012 into it manually.

VHD on IDE
0:59-1:12
If you have your VHDs on an IDE Controller, which is what the wizard is going to do, to add a new VHD to an IDE Controller requires you to power the VM off just like a real machine. I need to add an IDE Drive, I've got to turn it off.

VHD on SCSI
1:13-1:22
If you put your VHD on a SCSI Controller, then they're Hot swappable, and they can be added without having to power down the virtual machine. When we create our VHDs, we're going to have some choices.

Fixed Size
1:23-1:36
One choice is to make fixed size VHD. This makes the file as large as the maximum size of the VHD. If you make a 120 GB VHD, it'll grab 120 GB on the hard drive. This is your best performance. The only problem is, it may waste space.

Dynamically Expanding
1:37-2:13
Dynamically Expanding just expands the file to the size of whatever data I put in it. I can save space there. Maybe I know I need 32 GB to install Server 2012. I'm never going to use more than 10 GB. I can do Dynamically Expanding. I start copying a lot of data in that file, it's going to require more processing resources because it has to keep making the file bigger and bigger and bigger and bigger. If you have a Dynamically Expanding VHD and you delete items, it does not reclaim the space for the items that you deleted. So you put 40 GB of data in this file, the file size is 40 GB. You deleted 20 GB of data, that file size is going to stay 40 GB.
In a minute, we'll see what we can do about that.

Managing VHDs
2:14-2:52
Now, I can go through and I can perform maintenance on VHDs. I can Expand them. I can Shrink them. If I have a Dynamically Expanding VHD and I've deleted data from it, the way to reclaim that space is to Compact it. I can also Convert VHDs, so I can go back and forth between VHD and VHDX. I can go from Fixed to Dynamically Expanding, back and forth there. The thing to note is, when you convert a VHD, it makes a new file. That new file is going to have to be copied over the old file in order to get your VM booted up again.
Do not edit a disk for a virtual machine that has snapshots or replication enabled, or is associated with a chain of differencing disks.

Replication
2:53-3:12
Replication just means that the virtual machine is replicated to another Hyper-V server. It's a new feature of Hyper-V with 2012, that we can replicate VMs without having to install any clustering services. So that's the scoop on the new virtual hard drives available with Hyper-V on Server 2012, and some of the politics or the decisions that you'll make when you create your VHDs.

3.2.2 Creating Virtual Hard Disks

Creating Virtual Hard Disks
0:00-0:13
In this video, we're going to work with virtual hard disks inside of Hyper-V. I want to go into my Hyper-V Manager Console and go up to Tools, Hyper-V manager, and I am going to work with this Windows 8 virtual machine so that we can take a look.

Settings
0:14-0:26
We're going to right click "Go into Settings", and I can see this is my current virtual hard drive. If I need to make any changes to it I can do that here, but you can only do it when the machine is off, so if it were running it probably wouldn't let me do anything here.

Edit
0:27-0:31
Edit is how we can manage an existing virtual hard drive, and I have some options.

Compact
0:32-0:56
If it's Dynamically Expanding, which means the hard drive is only going to be as big as the data that I put inside it, that hard drive will grow as I add data. When I delete data, Hyper-V doesn't shrink the hard drive. Let's say I put 40 GB of data in this. The actual file is 40 GB big. I delete 20 GB of data. The file is going to stay 40 GB until I Compact it, and then that will actually shrink the file down.

Convert
0:57-1:27
I can Convert it. If it's currently VHD or VHDX, if it's Fixed or Dynamic, I can make changes. When you convert, it makes a completely different file, so after you convert, you're going to need to go back and copy the new file to the same location, same name as the old file. It doesn't matter what direction you're going from, it always makes a new file. I can go through and I'll get to Choose the Disk Format, Disk Type, Configure it ... I pick whatever I want.

Expand
1:28-1:34
If I need to add space to the VHD, I can expand it, and then I would be able to go in and specify how much space. If I want to make a new VHD, I can use the New button.

Make a New VHD
1:35-1:44
Now, I'm actually going to do it a little bit differently. I'm not going to do it in this window.

Inspect
1:45-1:54
Before we go, you can also Inspect the VHD and it will tell you the Current Size, Maximum Size, where it's located, File Name, and all the information you might need to know.

Make a New VHD Continued
1:55-2:22
I can make a new VHD right from this window, but I'm going to take you into another spot, because very often if I'm in this window, right now I'm on an IDE Controller. In a physical computer you can only work with IDE Drives when the computer is off. Like I said, if this computer were running, I wouldn't be able to do anything in here. If you have a physical SCSI Drive, on the other hand, those are Hot swappable and Hot addable. I'm going to show you how to create a new SCSI drive.

Pass Through Disk
2:23-2:53
Notice here I could actually use a real physical disk that's called a Pass Through Disk. The reason it's not available to me is that in Pass Through Disk the disk has to be in an offline state before I'll see it show up here. I would have to have another hard drive in the physical machine. I'd have to go into Disk Management, take it offline, and then I'd be able to select this. Technically, we call this a Pass Through Disk. If you see anything about Pass Through Disks, it means I'm actually using the hard drive directly and they've got to be offline before you can do it.

SCSI Controller
2:54-3:11
If this machine were running and I wanted to add a hard drive to it, I would click on SCSI Controller. You see I can add a hard drive, and then I would click Add. Now I can go through -- you can see it's the same as the IDE -- I'll hit New, and this is my New Virtual Hard Disk Wizard.

VHD or VHDX
3:12-3:41
The first thing that I need to choose is, is it going to be a new VHD or VHDX file? VHD supports up to 2040 GB in size. If I need it to be bigger than that, I need to go VHDX, which will give me up to 4 TB. It's also supposed to be a little bit more resilient to power failures.
The big thing with VHDX is, we're getting set up for newer and bigger virtual hard drives. I'm going to leave mine VHDX. That Convert in the Edit would let me swap between these two formats.
Disk Type
3:42-3:44
Now I choose my disk type.
Fixed Size/Dynamically Expanding
3:45-5:11
With a Fixed disk, it makes the VHD the size that you specify right away. If I say "this is 120 GB hard drive", it's going to grab 120 GB from my physical hard drive. The two good things about that are this: first of all, because I've already grabbed that space, there's no way I can create VHDs bigger than I have physical space available.
That's the only problem with Dynamically Expanding, is, I might actually get to a point where all the files dynamically expand to their maximum, and if I've got more VHDs configured than I have physical hard drive space, it could actually crash my host server. The other good advantage is, Fixed size is going to give you your best performance, because if I'm adding data to a Dynamically Expanding hard drive, the processor on the host machine has to keep making that file a little bigger, a little bigger, a little bigger, and a little bigger.
If performance is an issue, Fixed size is going to give me better performance. The only disadvantage to Fixed size is, I might be wasting hard drive space. Maybe I want to go through and make a bunch of 120 GB virtual hard drives, but I know for a fact I'm never going to put more than 10 GB into any of those machines. Well, Windows Server 2012 requires 32 GB free on the hard drive. Windows 8, I believe it's the same amount, so I can't just make it 10 GB Fixed size, because Windows won't install. If I know I'm only going to be using 10 GB, I don't want to waste 22 GB of space; then I can do Dynamically Expanding.
Differencing
5:12-5:21
Differencing disks are a little different, but basically what they're used for is to save space. We're going to pick Dynamically Expanding and hit Next.
After Choosing Disk Type
5:22-5:48
Now I go through, I give it a name, make sure it's stored in a location where you're going to have enough hard drive space for whatever you're going to do with this machine, and then I'll hit Next. I can set up the size of the virtual hard drive. This is new with Server 2012, where I can actually say, "Not only am I going to make this VHD, but I'm going to grab a copy of the contents of another VHD". That's pretty cool. I could also use it to grab the contents of a physical disk.
Virtualizing a Machine
5:49-6:06
If you're trying to virtualize a machine, you have a physical hard drive and you're trying to virtualize that machine, I can click this bubble, specify the disk that it's going to copy, and it will automatically do that for me without my having to boot up operating systems and that type of thing. We're just going to do a blank virtual disk, hit Next, and then hit Finish.

3.2.3 Differencing Disks

Differencing Disks
0:00-0:21
In this video, we're going to talk about Differencing Disks and Pass-Through Disks. Differencing Disks allow us to save space. They might require more processing resources for the virtual machine, so they're not really recommended in a production environment. The Differencing Disk must be based on the same format as the parent, which is .VHD or .VHDX, and then to merge a Differencing Disk with a parent, edit the Differencing Disk.

How it Saves Space
0:22-1:55
Let's take a look at how these things save space. Let's suppose that I have a really terrible hard drive on my computer and I've only got 20 GB overall, I need to install three servers. For the sake of argument, let's just say each Window Server 2012 requires 10 GB on the hard drive. Well, now I have a math problem. I need 30 GB of space, but I only have 20 GB. That's why we would use a Differencing Disk. Here's what I would do. I might make a 10 GB Dynamically Expanding VHD. I'm going to lie and tell it it's really 32 or 40 GB; but I'll install Windows Server 2012, it takes up 10 GB, that's how big my file is, let's say it's Windows Server 2012 .VHDX, I'm going to make this VHD Read-Only, and I'm not going to associate a virtual machine with it. It's just going to be VHD hanging out on the hard drive. Now, when I create my three servers, I give each one of these servers a Differencing Disk based off that original disk.
My DNS Server has a Differencing VHD that's just got the difference between a base install of Windows Server 2012 and whatever it takes to get DNS running. It's probably a couple of hundred MB. I'll do the same thing with my DHCP Server. I'll just add DHCP in, and that's all that's kept on the Differencing Disk. So, that might be another couple of hundred MB. Here, I install Active Directory. Now, there I might run into some problems if my Active Directory Database gets really big, but maybe this is just a test environment. Again, it's the difference between the base install of 2012 and installing Active Directory, and maybe here I use 500 MB. Even with all these Differencing Disks, I'm well within my 20 GB of physical space on my physical hard drive.

Uses of Differencing Disks
1:56-2:16
Differencing Disks let me save space, but generally they're only used for test computers. I've seen them used a lot in training, where student computers don't have very big hard drives, but we want a lot of VMs for the students to play with. For a Production Server, I wouldn't advise this. You're better off making sure you get enough space and just creating one VHD per computer.

Pass-through Disks
2:17-2:45
Pass-through Disks let me have the VM directly use the physical hard drive. There's no VHD involved, the VM is mapped to a particular spot on the hard drive. The only thing that you really need to know about Pass-through Disks is that the hard drive has to be taken off line before it will be available to the VM as a Pass-through Disk. Each of the volumes on that Pass-through Disk will only be accessible by one virtual machine. That's how we work with Differencing Disks and Pass-through Disks in Hyper-V.

3.2.4 Snapshots

Snapshots
0:00-0:38
In this video, we're going to talk about snapshots for Hyper-V machines. The idea behind a snapshot is, I'm at a certain point, and I'm not sure if the changes I'm going to be making are changes that I'm going to want to keep. When I take a snapshot, everything I do to the virtual machine after I made the snapshot is saved in a separate file, very similar to a differencing disk. If I don't like the changes, I can go back to the snapshot. I can revert, get rid of the snapshot, and just go back, or I can even have multiple snapshots and apply them as necessary and sort of hop around and test different features until I decide on exactly what I want.
Let's look at some facts about snapshots.

Changes Saved to a Separate File
0:39-0:47
First of all, as we said, snapshots allow the changes to be saved into a separate file so that the virtual machine can be returned to the state it was in as of the snapshot.

Space Reclaimed When a Snapshots are Deleted
0:48-1:01
Now, new with server 2012, when you delete a snapshot, the space is reclaimed without having to shut down the virtual machine. If you have earlier operating systems like 2008, 2008 R2, you actually had to shut down the machine before you got your physical hard drive space back.

Paused Critical
1:02-1:24
If the status of the virtual machine is Paused-critical, that means that the host machine is out of space for the snapshots. That's not a good thing. You can redirect snapshots to another drive, either when you create the virtual machine or after the fact. Make sure when you delete snapshots you always do it from inside of Hyper-V. Don't just go out to the hard drive and delete the snapshot file. It's going to completely mess up your virtual machine.

Merging Snapshots
1:25-2:02
When you're sure that you like the changes, make sure you merge snapshots as soon as the virtual machine passes testing if this is a production server.
A couple reasons: Number one, large files might take a long time to merge. As you get further and further away from the time when you took the snapshot, the snapshot files will get bigger and bigger. Now you're going to have a delay when they merge. They do have support in Server 2012 for Live Merge, so we don't have to have the VM offline. Still, performance isn't going to be as good as when it's not merging. Working off a snapshot also slightly reduces the disk performance of the virtual machine, so it's in your interest to merge that snapshot, and you'll get better disk performance.

Deleting a Snapshot
2:03-2:11
When you delete a snapshot, it merges it with the snapshot above. It's also going to delete any snapshots below. So be careful and make sure you are sure about what you're doing.

Snapshots Safe for Domain Controllers
2:12-2:24
Now, a new feature with Windows Server 2012 supports safe use of snapshots for domain controllers. It prevents the domain controllers from replicating old information, and it also coincidentally supports cloning domain controller virtual machines.

Domain Controllers
2:25-3:01
When you're working with snapshots, there is always a potential to get things messed up. With domain controllers, if I go back to a snapshot, it might be, if that snapshot was taken yesterday, now it's going to roll back a whole bunch of stuff with replication. In the old days, it would start replicating out old data. New with Windows Server 2012, it has a mechanism for keeping track of whether it's reverted to a snapshot or it's just booting up. If it sees that you reverted the domain controller from a snapshot, it actually marks the old stuff with a separate number so that it's not going to overwrite newer material on the other domain controllers.
Because of that, we can also use it to clone domain controllers.

Reverting Machines and Computer Accounts

3:02-3:50

A big problem that we often have with reverting machines is with computer accounts. We see it all the time in classrooms where you revert, let's say, a Windows 7 or a Windows 8 machine. Now the machine's relationship with the domain becomes corrupt, and you tend to get a message, "The trust relationship between this computer and the primary domain has failed".

What's happening is, the computer account for that computer has a password which gets changed every 30 days. The password got changed on the domain controller, you roll back the work station, and now the work station thinks it's the old password, and then the relationship between that workstation and the domain is now corrupt. There is actually a group policy that you can set, if that's a problem for you, that will go ahead and make sure that the domain controllers keep the old passwords for a little while so that if you do have to revert the machine, that trust relationship won't be broken.

3.2.5 Managing Snapshots/Checkpoints

Creating Snapshots
0:00-0:04
In this video, we're going to take a look at how to manage snapshots for virtual machines.
Windows Server 2012
0:05-0:20
I need to go in and get into Hyper-V Manager, Tools, Hyper-V Manager, and we'll work with some snapshots on this Windows Server 2012 and, so that we can see the snapshots, what I'm going to do is play around with the hard drive.
Working With the Hard Drive
0:21-0:29
Right now, we can see we just have the standard Windows Server 2012 folders in that hard drive. I'm going to go ahead and make a snapshot.
Make a Snapshot
0:30-0:39
Right click, Snapshot. Once the computer finishes making a snapshot, it appears here.
Rename Snapshot
0:40-1:17
If I need to rename it, I can rename it so I know what's going on with that snapshot.
Now, every change I make to this server is going to be kept in a different file. Let's go in on our server, and we'll just go ahead and make a folder in there After Snapshot 1. Now, I'm going to go back into Hyper-V Manager, and I'm going to make another snapshot.
If I need to roll back to a snapshot, I've got a couple of choices.
Roll Back to a Snapshot
1:18-2:11
I can actually go through and Apply the snapshot. If I do that, it will leave this snapshot here in existence. If I right-click the machine, and I click Revert, that's going to reset everything back to the most recent snapshot. I prefer to work with Apply.
We're going to come in here, and we're going to Apply this snapshot. Notice, I can delete this snapshot. I can delete the entire snapshot sub-tree as well. If I want to apply it, I can take a new snapshot and then apply it, or I'm actually just going to go ahead and apply this. You can see in the background, shut the server down. It's taking it right back to exactly what it looked like when I made that first snapshot.
There we are. Just for fun, we're going to make another change, and I'm going to go through and make a snapshot.
Used for Testing
2:12-3:36
I can go ahead and rename this snapshot. You wouldn't do this with a production server, but it's great for testing things, so that I can go through.
Maybe my base server is just an Install Server 2012, and then I want to install DHCP, and then I roll back. No DHCP, but I make a snapshot so that I can go back if I need to. I don't have to redo it, because I can apply any of these snapshots, so I can go here and say, Okay. I want to apply this snapshot, and it's going to take us back to that Folder 1. You can see there, After Snapshot 1. I can do the same thing with Snapshot 2, and now I'm back to After Snapshot 2.
Maybe I try a few different things, and I see what works. And I say, this is what I like. I like this one. I'm going to delete this snapshot, and I get rid of this one. Now, my base server has this, After Snapshot 2.
With Server 2012 R2, they've changed the name of snapshots to checkpoints and for the most part they're managed exactly the same as they were in 2012 except for the new name.
Server 2012 R2
3:37-3:48
If I wanted to go ahead I could right click and create a Checkpoint.

Create a Checkpoint

3:49-4:21

Just like a SnapShot, this is really a point-in-time picture of what the machine looks like in terms of RAM, in terms of what's on the hard drive.

There's a lot of different files involved with the Checkpoint. One of the things you might want to be aware of is it does take a snapshot of what's in RAM, so if you're looking to reduce the size of the checkpoints, you want to go ahead and turn the virtual machine off, that way you won't have to have the file that takes a snapshot of what's going on in RAM.

Apply or Revert a Checkpoint

4:22-4:47

Just like with snapshots, checkpoints we have the ability to apply them, which means we can apply that checkpoint and then build from there or if we just want to roll everything back, we can right click and go to Revert which will take everything back to exactly the way it was at that checkpoint and get rid of any checkpoints below.

If we delete the checkpoint, all the changes will be merged with the virtual machine.

Manage Checkpoints in PowerShell

4:48-6:15

You can go ahead and manage the checkpoints from inside of PowerShell.

Let's go into PowerShell and take a look at some of the commands that we can use. I can use a 'get-vm' and I give the name of the VM and I can see all of the information about it. If I'd like to create a checkpoint for that VM I can add the pipe and then 'checkpoint-vm' and you can see it's creating a checkpoint.

We go back into Hyper-V Manager, you can see immediately that the checkpoint has been created. If I want to go in and list all of the checkpoints I can do that as well--but for that we actually use that old word, snapshot. So I can go in and use my get-vm add the pipe and then get-vmsnapshot. It lists all of the snapshots for that virtual machine. So be aware with PowerShell, they haven't changed all the names of the commands, it might be checkpoint, it might be snapshot.

You can manage anything in Hyper-V from PowerShell. Anything from creating virtual machines through creating SnapShots, etc. But those are just some of the commands that we can use to work with snapshots from inside of PowerShell.

Other than that, besides the change from the word snapshot to checkpoint, there really is no difference with Windows Server 2012 R2 versus Windows Server 2012.

3.2.6 Virtual Machine Storage Facts

When creating a new Virtual Machine (VM), the operating system and other contents of the VM are stored in a VHD file.

Keep in mind the following when creating a new VM:

- When you create a VHD file during the process of creating a new VM, the VHD file will be a dynamically expanding disk.
- To use a fixed sized VHD with a VM, create the VHD before you create the VM.
- You have the option to copy the contents of an existing VHD or an existing physical disk to a new VHD.
- VHDs on IDE require the VM to be powered off in order to make changes.
- VHDs on SCSI are hot swappable, meaning that you can make changes to the VHD while it is running.

The following table describes management tasks you can perform on VHDs:

Action	Description
Expand	Expand allows you to increase the maximum size available in a VHD file.
Shrink	Shrink allows you to reduce the maximum size available in a VHD file.
Compact	Compact allows you to reduce the size of a dynamically expanding VHD. Dynamically expanding VHD files: • Increase in size as files are added. • Do *not* automatically reduce in size when files are deleted. The storage capacity of the VHD remains unchanged when you compact the drive.
Convert	Convert allows you to: • Change the format of a file from a VHD to VHDX, or change the format from VHDX to VHD. • Change the type from fixed to dynamically expanding or from dynamically expanding to fixed. When you convert a VHD file, it makes a new copy of the file. For the changes to take effect, copy the new file to the same location as the VHD file you changed.

Do not edit a disk with snapshots, replication enabled, or differencing disks.

Keep in mind the following regarding Hyper-V in Windows Server 2012:

- You can replicate a VM to a VM hosted on another Windows Server 2012 with the Hyper-V role installed.
- You cannot edit a running VHD.
- You can clone VM domain controllers.

The following table describes features you can use with VMs:

Feature	Description
Pass-through disk	When you connect a physical hard disk to a VM, the hard disk is referred to as a *pass-through disk*. • The physical hard drive must be on the virtualization server or it can be a network-attached disk. • The pass-through disk must be offline so the VM can have exclusive access to it. • The Disk Management snap-in cannot be used to add a pass-through disk to a VM.
Differencing disk	A *differencing disk* is a VHD associated with another disk and contains only changes to the associated disk. • The differencing disk is referred to as the *child disk*; the disk it is associated with is the *parent disk*. • The parent disk remains unchanged. The child disk contains the changes to the parent disk. • Differencing disks save space but require more processing resources. • A differencing disk must be the same format (VHD or VHDX) as the parent. • A differencing disk can be merged with the parent using **Edit** in Hyper-V Manager. • Differencing disks are not recommended for use in a production environment.
Snapshot/Checkpoint	A *snapshot*, also known as a *checkpoint*, can be used to restore a virtual machine to a previous state. • You can apply multiple snapshots to get to an exact state of the VM. • Space is reclaimed when you delete a snapshot. • In Windows Server 2012, a **Paused-critical** status indicates that the host machine is out of space for the snapshots and has been paused to allow the snapshot to be redirected to another drive. • Snapshots should be deleted through the Hyper-V Manager console.

- Merge a snapshot as soon as you determine the snapshot is the state to which you want to restore the VM.
    - Large snapshots take time to merge.
    - Windows Server 2012 allows you to merge a snapshot to a running VM.
    - Merging improves disk performance.
- When you delete a snapshot, the snapshot is merged with the preceding snapshot and deletes subsequent snapshots.
- In Windows Server 2012, a domain controller is prevented from replicating old information when a snapshot is merged.

In Windows Server 2012 R2, the **Checkpoint-VM** cmdlet is used to create a checkpoint on a VM. Keep in mind the following when creating a checkpoint:

- The **Checkpoint-VM** cmdlet lets you create the checkpoint while the virtual machine is running.
- The **-AsJob** parameter processes the cmdlet as a background job.
- The **-PassThru** parameter pipes the checkpoint.
- The PowerShell commands **Export-VM** and **Export-VMSnapshot** export a checkpoint.

3.3 Virtual Networks

As you study this section, answer the following questions:

- What is the difference between an *internal* virtual network and a *private* virtual network?
- Which virtual network types do not allow virtual machines to communicate with other hosts on the physical network?
- When would you need to use a legacy virtual network adapter?
- When setting up a virtual network, what is the recommended number of network cards for the physical system and how should they be configured?
- How can a virtual machine be configured to participate on a VLAN?

After finishing this section, you should be able to complete the following task:

- Create a virtual network.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 2.0 Hyper-V.
  - Manage Virtual Machines with Hyper-V Manager
    - Create Virtual Network and Settings

This section covers the following 70-410 exam objectives:

- 302. Create and configure virtual machine storage.
  This objective may include but is not limited to:
  - Implement a virtual Fibre Channel adapter
  - Configure storage quality of service (QoS)
- 303. Create and configure virtual networks.
  This objective may include but is not limited to:
  - Implement Hyper-V Network Virtualization
  - Configure Hyper-V virtual switches
  - Optimize network performance
  - Configure MAC addresses
  - Configure network isolation
  - Configure synthetic and legacy virtual network adapters
  - Configure NIC teaming in virtual machines

3.3.1 Virtual Network Adapters

Virtual Network Adapters
0:00-0:06
In this video, we're going to talk about virtual network adapters in Hyper-V. There are a few types of network adapters that you can add.

Synthetic (Network Adapter)
0:07-0:43
Traditionally, we've had Synthetic and Legacy Network Adapters. Now, you're not going to see the term Synthetic inside the operating system, it just says Network Adapter. But if it just says, Plain Network Adapter, technically Microsoft can call that a Synthetic Network Adapter. That's what they're talking about. These have better performance. They do require the integration services to be installed on the Guest Operating System; that's done by default. There shouldn't be any problems there. This I would pay attention to. These are the network adaptors that allow Bandwidth Management, and we'll talk about that in a minute, but you should definitely know that. We're not going to see Bandwidth Management if we install a Legacy Network Adapter.

Legacy Adapter
0:44-0:47
Now, Legacy is a politically correct term for old junk. Why would we want this?

PXE
0:48-1:12
Bizarrely enough, the Legacy Adapters are the only ones that support PXE booting a virtual machine, and that means booting the virtual machine off the network card. Why would we want to do that? Well, maybe we're trying to image the virtual machine with WDS. You would add a Legacy Network Adapter, make sure you're going in and adjust the boot order, boot it up, connect the WDS, pull down the operating system, and then when everything's said and done, get rid of the Legacy Network Adapter. That's what Microsoft recommends.

VLANs
1:13-1:21
In our Network Adapters, we can actually go ahead and specify VLANs. What VLANs do is separate traffic on one switch into different broadcast domains.

Bandwidth Management
1:22-1:33
They're also going to give us the ability for Bandwidth Management. I can set the minimum and maximum bandwidth that this particular network adapter is allowed to consume on the physical network adapter. If I set it to 0, that means it's unrestricted.

MAC Addresses
1:34-1:44
The other thing I can do with my Virtual Network Adapter is set a particular MAC Address. By default, it's set to Dynamic, and the VM will think up a MAC Address, but if you need a particular MAC Address, you can set that manually.

MAC Spoofing
1:45-2:26
They also have a setting to allow or disallow MAC Spoofing. If MAC Spoofing is not enabled, essentially what it does is protect you from malicious virtual machines. What happens is this, the virtual machine boots up and the switch says, "Oh, that's that MAC that goes to that virtual port", and that will become the only combination for that port. The VM won't be able to change its MAC Address. That port won't accept any other MAC Addresses. I can't boot up another VM, put in the same MAC Address as the first VM, and start sending out malicious packets. If you were looking to change this, you'd go into the Network Adapter, Properties, open that up, go into Advanced Features. Basically it just says whether or not we're going to let it change the source MAC Address for outgoing packets.

Fiber Channel Adapter
2:27-2:47
Those are some of the types of Network adapters we'll see.
There's one that's not on there. It would be a Fiber Channel Adapter, but that's only used if you have a Fiber Channel Host Bus Adapter in the host machine. Fiber Channel is a network for storage area networks. Really, we've gone through the Network adapters that I would expect you'd be working with in Hyper-V, and at least you know what each one is used for.

3.3.2 Adding Virtual Network Adapters

Adding Virtual Network Adapters
0:00-0:04
In this video, we're going to see how to add a virtual network adapter to a Hyper-V machine.
Windows Server 2012
0:05-0:13
I need to get into Hyper-V Manager, then to go Tools > Hyper-V Manager. I can only add a network adapter to a machine that's Off, so we're going to use this Windows 8 virtual machine.
Add a Network Adapter
0:14-0:42
Right-click, go into 'Settings'. Then here on the 'Add Hardware' page, I've got different types of network adapters that I can add. The one that says 'Network adapter' is sometimes referred to as a synthetic network adapter. It's the standard network adapter for Hyper-V. Legacy network adapter ' Legacy is a politically-correct term that means 'old junk.'
PXE Boot Support
0:43-1:47
I don't know why here, but the regular network adapter does not support PXE boot. So if you want to boot a virtual machine off its network adapter, let's say to connect it up to a WDS server and pull down an image, you have to add a Legacy network adapter to that virtual machine. They actually recommend you add it, connect up, pull down the operating system, and then later on, get rid of it instead of just leaving it inside the server. A fibre channel adapter would be added to connect to a fibre channel SAN. We would need a physical fibre channel host bus adapter in the host machine itself. I don't have one of those. I'm not going to add it. I'm just going to add a regular network adapter.
You can see, we can set up what virtual switch it's going to be connected to. I can connect it to a VLAN if I need to. It could be a VLAN within Hyper-V or VLAN on my network. I'm using VLANs on my network-- make sure you put in the identifier-- otherwise, you can just use them on the Hyper-V machine, and I can also set up bandwidth management. Just so you can see, if I add a legacy network adapter, it's really no different, other than the fact that I can't control the bandwidth, but again it will support PXE boot.
Windows Server 2012 R2
1:48-2:43
It's really not any different with Windows Server 2012 R2 versus Windows Server 2012.
One of the things you want to make sure that you do is provide multiple network cards for the virtual machine. To add a network card, we're going to need to shut it down. I'm going to connect to my virtual machine, sign in, and then shut down the virtual machine. Once the virtual machine is turned off, we're going to go ahead and add a network card. I'm going to go into 'Settings'. You can see I have one network card. I'm just going to add another one. I can connect it to the same network or a different network depending on what type of network team I want.
Enable NIC Teaming
2:44-3:54
The important thing if you're going to network team inside of the virtual machines is to go into the Advanced Features and enable NIC teaming on both of the network cards. You can see I'm going to enable this network card to be part of a team in the guest operating system. I'm going to do the same for my original network card. Once I turn on the settings of the virtual machine, I can reboot my server, and I'll be able to enable NIC teaming inside of the guest operating system. Now that I'm logged in to my virtual machine, for network teaming, we're going to go over to 'Local Server', and then we'll go ahead and enable network teaming.
We'll go ahead and create a New Team, Team1, add the network adapters. I can set up any additional properties that I need. If I want one as a standby or all of them active I can set that up. Then I click 'OK'. Really, no difference between network teaming in 2012 or 2012 R2.
Summary
3:55-4:04
Just make sure that you go into the properties of the virtual network cards that enable NIC teaming. Otherwise, it's not going to work.

3.3.3 Virtual Switches

Virtual Switches
0:00-0:11
In this video, we're going to talk about virtual networks in Hyper-V. There are really four types of virtual networks that you can create: External, Internal, Private, and then Fibre Channel.

External Network
0:12-1:06
First, we're going to start with External. Let's take a look at some facts about external networks. In order for my virtual machines to communicate, they have to be connected to a virtual switch. External virtual switches allow the VMs to communicate on the physical network.
Now, when we go in and we specify that in the properties of the adaptor, we're going to see a setting that says "Allow management operating system to share this adaptor". Here's what Microsoft recommends: Install two network cards on the Hyper-V Server, attach one to an external virtual network, specify that for the virtual machines, but uncheck this option so that the other network card is the only one available to the host. So, we're isolating host traffic from VM traffic.
When we set this up, by default, inside of the VM you're going to see a virtual network card. Even though we're communicating on the physical network, it's still going through the hypervisor, going through the software for the virtual switch.

SR-IOV
1:07-1:59
New with Windows Server 2012, there's an option for SR-IOV. Here's what this does: It allows the guest operating system in the VM to use the network card directly, without going through the hypervisor. If I go in the guest operating system, I'm going to see the actual network card that's installed in the physical box.
This is faster than the traditional external network. Here's the only sort of caveat to that: It has to be set when you're creating the network. When you create your external network, that's the time you want to check SR-IOV. If you don't check it then, you can't add it in after the fact. It's not a big deal, but it's something that you might want to know. It's a type of trivia that might show up somewhere. Again, it's recommended to use one adaptor for the host and a separate adaptor for external networks.
This SR-IOV is really cool because I can use the actual driver. It may actually give me some more functionality inside of the virtual machine than I would have without it.

Internal Network
2:00-2:14
If I specify an Internal Virtual Switch, that means that the guest operating system in the VM can access other virtual machines and the host, but it would have to go through the host to get to the real network. So it's not going to be out on the production network. It's just talking to other VMs and the Hyper-V host.

Private Network
2:15-2:26
If it's Private, that means that the VMs can only communicate with other VMs on the same switch, and that's generally used in testing.
You can also set a VLAN for the switch, or you can set the VLAN on the adaptor.

Definition of a VLAN
2:27-3:37
Let's take a minute and talk about what VLANs are, just so I make sure you understand. Here I just have an eight-port switch. By default, whatever traffic comes in to one of these ports, the switch will look at the MAC address and send it out to whatever port it's supposed to go out. If a broadcast comes in, it's going to go to all of the other seven ports. Now, let's assume you have security issues, or maybe there are just too many broadcasts. What you'd like to do is segment this switch down and have it operate as if it were two separate physical switches. That's pretty much what VLANs do.
So I could come in; say all right. These four ports are VLAN 2. These four ports are VLAN 3. Once I segment my ports into a VLAN, essentially the switch treats them as if they're two separate switches, or however many VLANs I make. That means when broadcasts come in, they only go out to the other ports in that VLAN. If I wanted somebody from VLAN 2 to talk to somebody from VLAN 3, I would actually need to set up a router between them, because as far as they're concerned, they're completely separate physical networks. If you have VLANs set on your physical switch, you want to set them in Hyper-V. If you just want to create VLANs in the Hyper-V switch, you can do that as well.

Virtual Fibre Channel SAN

3:38-4:17

The other type of switch that we can make would be a Virtual Fibre Channel SAN switch. There's a SAN Manager option in Hyper-V. This allows your VM to connect to a Fibre Channel SAN or Storage Area Network. In order to do this, you've got to have a Fibre Channel SAN. The host has to be connected to a physical Fibre Channel SAN, and the host has to have at least one Fibre Channel Host Bus Adaptor installed. You can actually go in and create up to four Virtual Fibre Channel adaptors in the virtual machine, and you can even associate each one with a different Fibre Channel Virtual SAN.

Again, you're not going to be working with this unless you have a Fibre Channel SAN, and then you can go ahead and dig into it.

Summary

4:18-4:43

In order for your virtual machines to use their virtual network adaptors, they have to be connected to a virtual switch. External switches let my machines access the physical network. Internal let them talk to other VMs and the host. Private isolates communication just among the VMs. Again, if I'm looking for better performance and my network adaptor supports it, I can turn on SR-IOV, and then the virtual machine will access the network card directly without going through the hypervisor.

3.3.4 Creating Virtual Switches

Creating Virtual Switches
0:00-0:18
In this video, we're going to see how to create virtual networks in Hyper-V. I need to get into Hyper-V Manager. We're going to go up to Tools, Hyper-V Manager. In my Virtual Switch Manager, I have three types of networks that I can create: External, Internal, and Private.

External
0:19-0:25
External allows the virtual machines to actually get access to the real network card in the server.

Internal
0:26-0:42
With an Internal network, the guest machine can talk to other virtual machines, so the virtual machines that are connected to that network can all talk together, and they can communicate with the host machine. They can go through the host machine to get to the internet if they need to, but they can't access the physical network directly.

Private
0:43-0:49
Private just allows the virtual machines to talk to each other. They cannot talk to the host. They cannot talk to the external network.

Creating an External Network
0:50-0:58
I'm going to make an external network because there's a few things that we can do with an external network that I can't do with the other two, but I can switch to the other two at any time if I want.

Naming the Virtual Network
0:59-1:10
I'm going to give my network a name, External Network, and you can see right now I can choose if I have multiple network cards which network card this network is going to map to.

Microsoft Recommendations
1:11-2:01
What Microsoft recommends is this: If you are going to have virtual machines using one of the network cards, go ahead and install two like we've done here. Set one for the VMs, but do not allow the management operating system to share this network adapter. That means that all the traffic from the host machine is going to go through the other adapter. This adapter will no longer be accessible to the host. That way, I can separate virtual machine traffic from host traffic.
If I just have one network adapter, you need to allow it to share it. Otherwise, the host operating system isn't going to have any network connectivity at all. This checkbox is only available when you're first creating the external switch. If you look down at the bottom, it says, SR-IOV can only be configured when the switch is created. Once I click off this page, I can't enable that.

VM's
2:02-2:25
Now, an external network allows the virtual machines to directly use the network adapter, so they can contact the physical network. By default, when I go into the guest operating system and I open up Device Manager, I'm just going to see a generic network card, because it uses a virtual network card. With that virtual network card, it's actually traveling through the hypervisor and then out the real network card.

Enable Single-Root I/O Virtualization
2:26-2:56
If I Enable single-root I/O virtualization, what that means is this: Now, when I go in my virtual machines, it will actually be able to use the driver for that physical network card, so I won't see the virtual network card driver. I'll see the actual network card in there, and it will go straight to the network card instead of through the hypervisor and then out to the network card. It should be slightly better performance. You just have to have a network card that supports this. If you don't see this checkbox here, it might mean that your network card doesn't support it.

Enable Virtual LAN Identification

2:57-3:24

I can also go through and Enable a virtual LAN ID, VLAN ID. If the physical switch requires a VLAN ID, you should turn that on. Once I get my network, I can click OK, or I can click Apply, and leave the box open. I should make as many networks as I need, and if I need machines to be separated, I can put them on different virtual networks, but they're going to have to be connected to some kind of a virtual network in order to have network connectivity inside the guest operating system.

That's how we manage virtual networks with Hyper-V.

3.3.5 Virtual Network Facts

A virtual network functions in much the same way as a physical network: two or more computers are linked together to share resources. In a virtual network, one or more VMs are linked together or to an external network. Virtual components can be used with VMs to create virtual networks.

The following table identifies types of virtual networks:

Network Type	Description
External	In an *external* network, VMs bind to the physical NIC, which allows them to access the physical network. Use this option to allow the VM to communicate with the host operating system, other VMs running on the system, and other physical network devices.
Internal	In an *internal* network, VMs can communicate with one another and with the host operating system, but cannot access the physical network. This configuration is typically used to build a test network where you connect to the VMs through the management operating system.
Private network	In a *private* network, VMs can communicate with each other but cannot communicate with the host operating system or access the physical network. This network type is optimal when a VM needs to have a degree of isolation.
Fibre Channel	In a *Fibre Channel* network, VMs can connect to existing Fibre Channel storage arrays.

If you do not set up a network, the VM cannot communicate with any other physical or virtual machine. It may be best to not have any network setup for your VM if it is used for testing or performing other tasks in which isolation is beneficial.

You can use virtual network adapters to:

- Manage bandwidth settings for the physical adapter:
  - Minimum bandwidth
  - Maximum bandwidth
  - Unrestricted bandwidth
- Manage MAC addresses for the virtual adapter:
  - Disable MAC spoofing to protect from malicious VMs.
  - Enable MAC address spoofing to allow a VM to change its source MAC address for outgoing packets or to support failover in NIC teaming.

Use Hyper-V Manager to create and manage virtual adapters. The VM must be offline in order to create, configure, or manage the virtual network adapter.

The following table describes which types of virtual network adapters are supported in generation 1 and 2 VMs:

VM Generation	Virtual Network Adapter Types
Generation 1	Generation 1 virtual machines can use the following types of virtual network adapters: • Synthetic network adapters, also referred to as network adapters. Synthetic network adapters: Provide better performance than legacy adapters. Require Integration Services. Allow bandwidth management. • Legacy adapters are older network adapters. Legacy adapters are: Required to PXE boot a VM. Used for imaging VMs with WDS. • Fibre Channel adapters are used with Fibre Channel Storage Area Networks (SANs).
Generation 2	Generation 2 virtual machines can use the following types of virtual network adapters: • Standard (synthetic) network adapters. In generation 2 virtual machines, legacy adapters are no longer used. Standard network adapters in generation 2 virtual machines now include the ability to perform a PXE network boot. • Fibre Channel adapters are used with Fibre Channel Storage Area Networks (SANs).

Similar to a physical switch, a virtual switch allows one or more VMs to transmit data to local or external network resources. Be aware of the following regarding virtual switches:

- Virtual switches created by Hyper-V are software-based.
- Virtual switches can have an unlimited number of ports. Ports can be added or removed dynamically.

The following table describes considerations for installing virtual switches:

Type	Description
External	External virtual switches allow a VM to communicate with the external network as well as other VMs. In Windows Server 2012, Hyper-V provides:

	• Single-root I/O Virtualization (SR-IOV) Allows the guest OS to use the physical network adapter directly Provides faster communication speeds than going through the hypervisor SR-IOV must be set up when you are creating the network. • An **Allow host OS to share the network adapter** option. Microsoft recommends for the physical system to have two network cards installed and configured as follows: Attach one to an external virtual network and specify it for the VM. Uncheck the **Allow host OS to share the network adapter** option.
VLAN	VLANs can be set on the virtual switch, the physical switch or on the virtual adapter. • VLANs can increase security and reduce broadcast traffic. • To create a VLAN on a virtual switch, segment ports on the switch.
Fibre Channel	Fibre Channel SAN switches: • Allow the VM to connect to a Fibre Channel SAN. • Require a physical Fibre Channel SAN and a Fibre Channel host bus adapter (HBA). • Support up to four virtual Fibre Channel adapters on a VM. • Allow you to associate each Fibre Channel adapter with a virtual SAN.

3.3.7 Network Optimization

Network Optimization
0:00-0:15
In this video, we're going to talk about virtual network optimization. We're going to look at a lot of different features. Your best bet is just make sure you know a basic description of each feature, and then you should be good to go. Most of the stuff is in the properties of the virtual network adapter.

Virtual Machine Queue (VQM)
0:16-0:46
The first one we'll talk about is the Virtual Machine Queue. This is just turned on, and there's no reason to turn it off. What it does is allow network traffic that comes into the virtual card to be distributed across multiple CPU cores. The physical network card must support this feature, and the virtual machine has to have more than one core allocated to it. It's a performance benefit. If we can spread out the network traffic, it just speeds things up a little bit. It's either supported physically or it's not. It's checked in the virtual machine by default. If you've got it on the physical card, you've got it in the virtual machine.

IPSec Task Offloading (IPSecTO)
0:47-1:32
The next one is IPSec Task Offloading, which they sometimes call IPSecTO. It allows the security association for IPSec to be transferred to the physical NIC for processing. Now, this is another one where the physical NIC must support this, and you can go in and set a maximum number of security associations that can be handed down to the NIC. Without having an entire lecture on IPSec, let's just say that IPSec is used to encrypt network traffic. Every time two computers establish an encrypted channel, they set up some session keys which are called Security Association (SAs). Instead of having the guest operating system inside the VM use those keys to encrypt traffic, we can pass that off to the network card, which will free up some of the resources inside the virtual machine. Again, this has got to be supported by the physical network card.

DHCP Guard
1:33-1:56
The DHCP Guard prevents this particular virtual machine from being used as a rogue DHCP Server. Essentially what it does is, it tells Hyper-V if there is any DHCP packets coming out of this virtual network card. Just drop them so that nobody can go in and set up DHCP inside the the virtual machine and then start messing up your network. Windows servers already have protection from rogue DHCP servers in the sense that they have to be authorized. This is just an extra level of protection.

Router Gaurd
1:57-2:08
We also have Router Guard, which pretty much does the same thing, except now we're preventing the virtual machine from becoming a rogue router. Rogue -- if you're not familiar with the term -- that's just anything that you didn't authorize on your network.

Port Mirroring
2:09-2:45
Inside of the properties here, we're also going to see port mirroring. Port mirroring copies traffic to another port. It's used for Packet Sniffers, Intrusion Detection Systems, Intrusion Protection Systems, and anything else where one computer needs to look at the network traffic of another computer. When you set it up, you specify the port as either the source or the destination.
The great thing about a switch is that it learns the MAC address of each computer that's connected to the port. When traffic comes in from one port, it's routed just to the port with the MAC address for the destination. That reduces collisions in our network and makes the network run faster.

Packet Sniffer
2:46-2:55
The problem is, if you want to use a packet sniffer to look at the traffic, you're only going to see the traffic destined for the machine running the packet sniffer, which is probably not what you want to look at.

Intrusion Detection Systems and Intrusion Protection Systems
2:56-3:09
Intrusion detection systems and intrusion protection systems also need to look at network traffic so they can detect when a hacker is making an attack on the network. Detection systems just detect attacks. Protection systems actually shut down the attack once it's detected.

NIC Teaming
3:10-3:24
We also have an option for NIC teaming, which allows the network card to be part of a team in the guest operating system. If you don't turn on the properties of the network adapter in Hyper-V, even though you make a team inside the guest, it's not going to work. If one fails, it's just going to drop the other one.
Port Access Control Lists (ACLs)
3:25-4:00
We also support for Port Access Control List. Any time we have an access control list it means we're controlling who has access. So these are rules that you can apply to a Hyper-V switch port. You can specify whether a packet is allowed or denied into or out of the virtual machine. We can specify it by local address or remote address, but we can't do both. We can specify it IPv4 address, IPv6, or even MAC addresses. In fact, I can set up a whole network. Nothing is allowed from the 10.000 network. I can set up a direction, have this affect inbound or outbound packets or both, and then say whether the packets will be allowed or denied.
Example PowerShell Command
4:01-4:22
Here's kind of an example PowerShell command where we're adding a VM network adapter ACL. I specify the name of the virtual machine here. Here we're specifying by MAC address, and it's a local MAC address, and we just put down the MAC address. The direction is Both, and the Action is to Allow. It's sort of like a rudimentary firewall built in to the virtual network switch.
Meter Port ACLs
4:23-4:31
We also can set up Meter Port ACLs, which will measure how much traffic is sent to or from the virtual network card, and a specific address range.
Example PowerShell Command
4:32-4:53
Here's sort of an example. I'd be Adding a VMNetworkAdapterACL. I specify the name of my VM. The RemoteIPAddress I'm talking about is anything in the 192.168.0.0 network with a /16 bit subnet mask. It's Outbound traffic, and I'm just going to record how much, presumably because I'm going to bill whoever is in this network for whenever the server is talking to them.
Private Virtual Local Area Network (PVLAN)
4:54-5:29
I can set up Private Virtual Local Area Networks, and I can set them up for promiscuous or isolated if I use network isolation and logically segmenting the traffic. What I'm trying to do is put this computer in a VLAN by itself, either for security or for performance.
So my PVLAN can be used to create an environment where VMs can only interact with the internet and will not have any visibility to the other virtual machine's network traffic. If I wanted every single computer to be like this, I could get them all into the same PVLAN in isolated mode. Then I'd really only need two VLAN IDs, a primary and a secondary for each server.
Example PowerShell Command
5:30-5:43
Here's an example of a command that would do that. It's Setting the Virtual MachineNetworkAdapterVlan. I specify the name of my virtual machine. Here in this case it's Isolated. Then I give it a Primary and SecondaryVlanId.
PVLAN Modes
5:44-5:48
When you're working with PVLANs, there are three modes that we can use.
Isolated
5:49-5:55
Isolated means it's only going to communicate with the promiscuous ports in the PVLAN, which is only going to be the port leading to the internet.
Promiscuous
5:56-5:58
Promiscuous ports communicate with everybody.
Community
5:59-6:06
Then we also have a setting that's called Community, which means it communicates with ports in the same community and any promiscuous ports in the VLAN.

Trunking

6:07-6:28

We also have support for Trunking. What trunking does is allow a virtual machine to see traffic from multiple VLANs. These are technologies that exist on regular switches. Microsoft is now building them into the virtual switches. It's all got to be done by PowerShell. A switch port receives traffic from all the VLAN's that you configure, and you set up an allowed VLAN list.

Example PowerShell Command

6:29-6:45

Here, I'm setting up my VirtualNetworkAdapterVlan. I give the virtual machine name. Here's it's just VM, but you put the name of the VM here. Trunk means I'm setting up trunking. Here's my AllowedVlanIdList, and it's allowing 1 through 100. The NativeVlanId for this is 10.

Summary

6:46-7:01

So those are some of the features that we can use to optimize our virtual networks. Most of them are in the properties of the adapter or done with PowerShell. Like I said, I don't think you should memorize a lot of PowerShell commands. Just have an idea of what each feature does, and you should be all set.

3.3.8 Optimizing Virtual Network Performance

Optimizing Virtual Network Performance
0:00-0:09
In this video, we're going to look at Optimizing Virtual Network Performance inside of Hyper-V. The first thing we want to do is go into Hyper-V Manager, and click Tools, Hyper-V Manager.

Settings
0:10-0:17
We're going to take a look at the Settings for this Windows 8 virtual machine, because not everything is available if the virtual machine it running.

Network Adapter
0:18-0:26
Now, I can go through with my Network Adapter and select which virtual network it's connected to. If I need to, I can Enable VLAN Identification.

Bandwidth Management
0:27-0:48
In terms of optimizing, the first place we see optimization is here with Enable bandwidth management. I actually can limit this network adapter to a minimum and a maximum bandwidth across the physical network adapter. What that does is prevent this virtual machine from taking over the network card, so that's the first place we can do optimization.
If I click the plus (+) next to my Network Adapter, we see two sub choices, and we're going to first look at Hardware Acceleration.

Hardware Acceleration
0:49-0:54

Virtual Machine Queue
0:55-1:41
By default, the Virtual Machine Queue is turned on, and basically what that does is this. If the physical machine and the physical network card support spreading out the network traffic across multiple cores; so maybe I have a quad core CPU, and the network card has to support it, the physical card has to support it. So I spread out all the network traffic that's coming in across all four of those processing cores, so that all four cores are working on network traffic. If I don't support this, then it means that just one of the cores is working on network traffic. It might be faster if VMQ is working. This just says, "Enable VMQ if the physical adapter supports it". If the physical adapter doesn't support it, this check box means nothing, because if I can't physically support it, I can't do it.

IPSec Task Offloading
1:42-2:24
IPSec Task Offloading also has to be supported by the physical network adapter. The guest operating system has to support it as well, but basically what it means is this: In IPSec, we're actually sending encrypted traffic, and when we exchange keys that we use to encrypt that traffic, it's called an SA -- Security Association. What the virtual machine can do is this. It can establish the IPSec communication, but then give that SA to the network card so the network card is doing the actual encryption, which takes the burden off the operating system. If the operating system is doing a lot of IPSec and it's slowing down the OS, if it's supported by the network card, I can actually just hand that off to the network card and free up some resources inside my virtual machine.

Single Root I/O Virtualization
2:25-2:43
Single Root I/O Virtualization allows me to get direct access to the network card. I need to have a physical adapter that supports this, and I also need to create an external network and turn this on, and you can only turn it on when you're creating the external network, not after the fact.

Advanced Features
2:44-2:46
We also have some advanced features

MAC Address
2:47-3:42
Right off the bat, up top, we can set the MAC address. By default it's set to Dynamic, so Hyper-V will make up a MAC address. If you need to set your own MAC address for whatever reason, you can set a Static MAC address. You can also click "Enable MAC address spoofing". If MAC address spoofing is not checked, when this virtual machine boots, it's going to give its MAC address to the virtual switch. At that point, the virtual switch will say, okay, on this port in the virtual switch, that's the only MAC address I'm going to allow. That prevents somebody taking over this machine,

spoofing another computer's MAC address and using that to attack the network. If I turn on MAC address spoofing, the virtual switch will allow any MAC address on that port, and potentially I could have a security issue. Now, if you do have a situation where you need to allow MAC address spoofing, it's there so you can turn it on, but unless you actually need that, don't turn it on.

DHCP Guard

3:43-4:06

DHCP Guard basically protects me from this virtual machine becoming an unauthorized DHCP Server. If I enabled DHCP Guard, even if I go in and install DHCP on the server, the virtual machine is going to drop all DHCP traffic coming from this virtual machine. It won't be able to be a DHCP Server, and that prevents it from being a rogue DHCP Server.

Router Gaurd

4:07-4:25

The same thing with Router Guard. Even if I install RRAS inside the virtual machine, I create a router out of the virtual machine, if I turn on Enable router advertisement guard, the virtual network adapter is going to drop all router advertisement and redirection messages from that machine. It just won't even let it out of the port.

Port Mirroring

4:26-5:41

Port mirroring basically allows the network traffic of one virtual machine to be monitored by another virtual machine. When I mirror a port, what it means is, everything coming out of Port A or coming into Port A is also going to be sent to Port B. Normally, switches will only send traffic to a port if there is the correct MAC address on that port, and that makes them work efficiently because it cuts down on collisions. The problem comes in if I need to work with some type of a software that requires port mirroring. Maybe I want to run a packet sniffer on a different computer to look at the traffic for this virtual machine, because I'm looking for some kind of a problem or I'm looking for particular traffic that's coming out of it. The monitoring machine is never going to see the traffic from this machine unless I turn on port mirroring.

It's also used if you have computers that are intrusion detection systems or intrusion prevention systems. If they're network IDS or IPS, they need to see all the traffic on the switch so that they can detect or prevent a network intrusion -- basically a hacking attack. When I set up port mirroring, I can specify whether this virtual network card is the destination or the source.

NIC Teaming

5:42-6:11

And then finally, NIC Teaming allows me to configure NIC Teaming inside the virtual machine. If I check this, it will actually allow me to set up a NIC Team inside the VM. If I don't check it, you can see right at the bottom there, if you create a team in the guest operating system, it's going to lose connectivity if one of the network adapters stops working, even though I have a team set up. The team inside the VM is not going to work unless I come in here and Enable NIC Teaming.

VLAN Isolation

6:12-6:35

Now, there are some optimization techniques that we can only perform in PowerShell. Just to make you aware of these, one of them is VLAN Isolation, where I can go through and set up a virtual machine to be isolated on its own VLAN, and what that does is protect it from anybody getting in there in terms of a network. It basically ends up a network of one.

Access Control Lists

6:36-6:45

I also can set up Access Control Lists on the port, so I can go through and say, well, I'm only going to allow traffic from certain networks back and forth between this virtual adapter, and I can turn on Resource Metering, which means I'm going to record network statistics or virtual machine statistics so that I can bill clients if I'm offering hosted Hyper-V services.

Resource Metering

6:46-6:58

Those are some of the things that we can do to optimize Hyper-V networks.

3.3.9 Optimizing Virtual Network Facts

Features available to optimize virtual networks are described in the following table:

Feature	Description
Virtual Machine Queue (VMQ)	The Virtual Machine Queue (VMQ) allows network traffic received on the virtual network adapter to be distributed across multiple CPU cores. • This feature is enabled by default. • The physical network adapter must support this feature. • The VM must have more than one core allocated to it.
IPSec Task Offload (IPSecTO)	IPSec Task Offload (IPSecTO) allows the security association (SA) for IPSec to be transferred to the physical network adapter for processing. • SAs are security keys for the encrypted session. • In IPSecTO, traffic encryption is performed by the physical network adapter. • The physical network adapter must support this feature. • You can specify a maximum number of SAs that can be transferred to the physical network adapter.
DHCP Guard	DHCP Guard prevents a VM from being used as a rogue DHCP server. When DHCP Guard is enabled, the hypervisor drops any DHCP packets coming from the network adapter.
Router Guard	Router Guard prevents the VM from being used as an unauthorized router.
Port Mirroring	Port Mirroring copies traffic to another port. Port Mirroring is useful when evaluating traffic on the network, such as: • Packet sniffers • Intrusion Detection Systems (IDS) • Intrusion Protection Systems (IPS) When you set up port mirroring, you specify the port as either the source or the destination.
NIC Teaming	NIC Teaming, also known as Load Balancing/Failover (LBFO), allows a network adapter to be part of a team in the guest OS. NIC teaming must be enabled on the network adapter in the hypervisor to work. Be aware of the following facts about NIC teaming:

- NIC teaming provides bandwidth aggregation.
- NIC teaming provides traffic failover if a network component fails.
- You cannot run NIC teaming on a Windows 8 or Windows 8.1 Hyper-V VM.
- You can use the NIC teaming User Interface and NIC teaming Windows PowerShell cmdlets on a Windows 8 or Windows 8.1 machine to manage teaming on Windows Server 2012 or Windows Server 2012 R2.

When configuring NIC teaming keep the following in mind:

- You can configure NIC teaming in one of the following ways:
  Switch-independent teaming allows adapters in a team to be connected to different switches. If the NIC team is being used for failover only and not bandwidth aggregation, the NIC team must be configured as switch-independent.
  Switch-dependent teaming requires adapters to be connected to the same switch. You can implement switch dependent teaming in one of the following two ways:
    - Generic or static teaming requires that the switch and the host identify the links in the team.
    - Link Aggregation Control Protocol (LACP) teaming uses LACP to dynamically set the links between the host and the switch.
- Windows Server 2012 R2 supports the following traffic load distribution algorithms:
  Hyper-V switch port algorithm directs traffic using the VM's MAC address or the port on the Hyper-V switch to which the VM is connected. This algorithm is generally effective unless the host has only a few VMs.
  Address hashing algorithm creates a hash value using the address of the packet. Packets with that hash value are assigned to an adapter. You can use Windows PowerShell cmdlets to specify how the hash is created:
    - Using source and destination TCP ports and source and destination IP addresses. This is the default creation method.
    - Using source and destination IP addresses only.
    - Using source and destination MAC addresses only.

  Dynamic algorithm combines the Hyper-V switch port and the address hashing algorithms.
    - Outbound loads use a hash of TCP ports and IP addresses, rebalancing loads in real time. This methods also breaks TCP flows at naturally occurring breaks known as *flowlets*.
    - Inbound loads are distributed as if the Hyper-V port mode was in use.

	• Windows Server 2012 R2 supports NIC teaming in a VM. Virtual NICs connected to more than one Hyper-V switch can remain connected even though a physical NIC disconnects from the switch. • Multiple team interfaces can separate inbound traffic by VLAN. • If a team is connected to a Hyper-V switch, segregate the VLAN using the Hyper-V switch, not the NIC teaming software. When using NIC Teaming, enable MAC spoofing to allow traffic from an alternate network adapter in the event of failover.
Port Access Control Lists (ACLs)	Port Access Control Lists (ACLs) establish rules applied to a Hyper-V switch port that determine if a packet is allowed into or out of a VM. Port ACL rules have the following components: • **Address** specifies a local address or a remote address using an IPv4, IPv6, or MAC address. • **Direction** indicates the traffic direction the rule applies to: inbound, outbound, or both. • **Action** allows or denies the traffic. For example, the following command allows the VM named vm206 to receive packets from and send packets to a local MAC address of 12-46-56-83-97-7C: **Add-VMNetworkAdapterAcl -VMName vm206 -LocalMacAddress 12-46-56-83-97-7C** **-Direction Both -Action Allow**
Meter Port ACLs	Meter Port ACLs allow you to measure virtual adapter traffic sent to or received from a specified address range. For example, use the following command to measure outbound traffic to the remote address 189.207.0.0/24 from the VM named vm206. **Add-VMNetworkAdapterAcl -VMName vm206 -RemoteIPAddress 189.207.0.0/24** **-Direction Outbound -Action Meter**
Private Virtual Local Area Network (PVLAN)	Private Virtual Local Area Network (PVLAN) allows you to configure switch ports to control VM communication. A PVLAN has three modes: • Isolated mode allows communication with only promiscuous ports in the PVLAN. • Promiscuous mode allows communication with all ports in the PVLAN. • Community mode allows communication with ports in the same community and any promiscuous ports in the PVLAN. Creating a PVLAN in isolated mode allows you to restrict the traffic of all VMs in the PVLAN using only two VLAN IDs: the primary and the secondary. The

	following command isolates a VM named vm260 using the Primary VLAN 15 and the Secondary VLAN 176. **Set-VMNetworkAdapterVlan -VMName vm260 -Isolated -PrimaryVlanId 15 -SecondaryVlanId 176**
Trunking	Trunking allows a machine to see traffic from multiple VLANs. When trunking is enabled, a VM's switch port receives traffic from all VLANs configured in an allowed VLAN list. For example, a command to allow the VM named vm110 to view traffic from VLANs 3 - 350 is shown below. In this example, traffic without a specified VLAN is handled as if it is from VLAN 322. **Set-VMNetworkAdapterVlan -VMName vm110 -Trunk -AllowedVlanIdList 3-350 -NativeVlanId 322**
Storage quality of service (QoS)	Storage quality of service (QoS) allows you to control the throughput of data to virtual disks. In Windows Server 2012 R2, you can: • Use Hyper-V Manager to configure the maximum and minimum Input/Output operations per second (IOPS) value for each hard disk in a VM. Normalized IOPS is used for settings. IOPS is measured in increments of 8 KB. If no maximum is set, the system defaults to zero. • Use the WMI interface or Windows PowerShell to control and query the IOPS values. • Monitor and manage the effect of a virtual disk's IOPS on other virtual disks in the VM. • Define thresholds and receive notifications when thresholds are not met. • Manage individual tenant I/O throughput in a multitenant environment. • Collect data throughput information for chargeback. Storage QoS is not available for shared virtual hard disks.

4.1 Active Directory

As you study this section, answer the following questions:

- What are the advantages of a client-server network model versus a workgroup model?
- What is the difference between a *tree* and a *forest*? How can you tell when a new domain starts a new tree?
- What is the function of transitive trusts in a forest?
- What is the function of the schema?
- How does Active Directory ensure that each domain controller has the most current information from other domain controllers?

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Manage Active Directory.

This section covers the following 70-410 exam objective:

- 500 Install and Administer Active Directory.

4.1.1 Active Directory Structure

Active Directory Structure
0:00-0:08
What we need to do in terms of active directory is get a feel for the overall structure, because when we promote a domain controller, we want to know where it's going to fit in, in the organization.

Four Levels of Organization
0:09-0:48
There are four levels of organization in active directory. At the very top level is the forest. Within the forest are trees. Within the trees are domains. They're made up of the domains, and then inside the domains, we create organizational units in order to structure our organization. All of these should mirror the logical structure of your organization. It doesn't have to worry about the physical structure until we get to sites. We're going to go through and take a look at some of the features at each of these levels, so that you can be aware of them when you are going through promoting your domain controllers and creating your active directory structure.

Forest
0:49-1:09
At the highest level, we have the forest. Microsoft doesn't really have a way to draw the forest. I kind of draw it like this. Sometimes, if you look at diagrams from Microsoft, it's difficult to tell whether something's in the same forest or not. At the highest level is the forest. It's not easy to tell whether domains are in the same forest. You just have to look around.

Forest Root Domain
1:10-1:22
The very first domain that you install into this forest -- and domains are usually drawn as triangles -- is called the Forest Root domain. If you see documentation that says, "This is the Forest root domain", it means it was the very first.

Trees
1:23-3:50
Within the forest are trees. Trees are domains that are in a parent-child relationship and share name space. If our forest root domain is Northsim.com, then this might be West.Northsim.com. I know that these two domains are in a tree because they share this name space. They both have Northsim.com. The .com doesn't count, so it's got to share a piece of the name. This is done when I install the domain. I identify whether it's a new domain, a new tree, a new forest, or all three of those.
Within the forest, you can have as many trees as you need to. That being said, keeping it simple is the very best policy you can follow, so I definitely do not want more than one forest for my organization. Usually a forest is the entire organization. The only time you should end up with more than one forest is if you have a merger of companies. The reason that we break things down into trees is really arbitrary. It might be because we want to make things easier for the users. Let's say I own a company, and my company has a section where we make cookies, and we sell cookies in the grocery store. But then, we decide to diversify our interests, and we buy another company that does tires. Well, the people in cookies.com are really not going to want to log into a domain named tires.com, so I might make two different trees, one cookies.com, one tires.com, but within the same forest. The child domains allow me to break down those pieces for subadministration.
If I have a situation where now the cookies part of my company becomes so big that it's got more stuff in it -- more users, more computers, more servers -- than is practical for replication between them, then I might want to say, Well, I'll have my main name "cookies.com", but then I'll start to break it down. Could be east and west. Could be sales and production. However that makes sense. But, I'm looking to separate out pieces of it, so the database in any one domain is not too big. That's really the only reason why we would make separate domains, either as a security situation where I want to break people into different domains because they shouldn't interact with each other for security purposes, or because that database for the one domain has become too big and I need to separate things out.

Using One Forest
3:51-4:07
The reason we want everything within one forest is because being together in a forest gets us some advantages. Take a look back at our diagram. The parent domain - this one here, which happens to also be the forest root domain - has a relationship with a child domain.

Trust
4:08-5:15
In order to allow users in one domain to use resources in another domain, that's done with a trust. Now, in the days before Active Directory, you used to have to make trusts manually. The fantastic thing about being in a forest together, is if we have multiple trees; so I'll call this westsim.com, these are different tree roots. There's a two-way trust between the tops of the trees, which means that all the users in northsim can use all the computers in westsim, and all the users in westsim can use all the computers in northsim. There's also a two-way trust between the parent and the child within the tree. Again, all the users in northsim can use all the computers in westsim, and vice versa. These trusts are transitive, which means that they're being a trust between the parent and child, and a trust between the tree roots. Effectively, any user in any domain can use any computer in any other domain, simply by virtue of the fact that we're in the same forest. Anything that's in a different forest we don't have a relationship with, so that's why we want one forest for the entire organization.

Two-Way Transitive Trust
5:16-5:32
The only other thing that we get in common by being in this forest, besides these two-way transitive trusts - this is benefits of a forest, two-way transitive. Transitive means, if A trusts B and B trusts C, A and C have a relationship. So A trusts B, B trusts C, then there's a relationship between A and C.

Schema
5:33-6:11
That's transitive. The second benefit of being in the same forest is that we all share a common schema. The schema is the template for Active Directory. It's made up of classes and properties. Basically what it does is define objects in Active Directory. Within the schema is a definition of what is a user. User has a user name. It has a full name. It has a logon name. If you make a change to the schema, every single user and every single domain in your forest is going to change. When you go through with other products that deal with the schema, like Exchange - you can have one Exchange organization per forest - it's going to modify the schema to make way for Exchange.

Summary of Benefits
6:12-6:29
Those are my two benefits of being in a forest together, is that there are two-way transitive trusts between the roots of the trees, and between the tops of the trees and the child domains in the trees, and that we all share a common schema. Again, we want just one forest per organization, if we can at all manage that.

Summary
6:30-6:57
But as you go through, depending on where your company is at, if you're just installing active directory for the first time, you're going to be creating a forest, a tree, and a domain all at once. If you have an existing forest, and you're installing a domain controller, you have to make a decision. Is this a new tree? Is this a new domain in an existing tree? Or, is this just another domain controller inside of a domain that already exists, which we call a replica domain controller.

4.1.2 Computer Roles

Computer Roles
0:00-0:32
Now we want to talk about the roles that computers can have within an organization. First of all, your organization either has a Workgroup, or it has an Active Directory structure domain. If it's a Workgroup, that means there's no centralized database. Each of the computers is going to manage its own security. Active Directory really is a centralized database; that's what it does. It allows users to log in once, and be authenticated, and then have that authentication be valid throughout the domain.

Workgroup
0:33-0:42
If I have a Workgroup, it doesn't have to be clients, it can be servers. Let's say I have two computers. Whether they're running Windows Server or they're running a client operating system makes no difference.

Security Accounts Manager (SAM)
0:43-2:12
Each of these computers has a database in it called the SAM, which stands for Security Accounts Manager.
You don't even have to know what that stands for, we just know that each of these has a database, this has its own SAM, this has its own SAM. If you're setting up this Workgroup, you're going to have to go ahead and make a Shad account at each one of these computers, because they handle their own security and you're going to have to synchronize my password on each computer. As long as you set me up identically on all these machines, everything will work great. Make a Shad account here, and a Shadow account there -- well I'm going to be very sad because every time I'm sitting at computer A and I go to use something on B, I'm going to have to re-put in my username and password.
The beauty of Active Directory is that we have a central database which is Active Directory. That database lives on a domain controller. Domain controller, as soon as I say domain controller, tells me it has a copy of Active Directory. When I go through and I sit down at my computer, instead of my computer handing my username and password to the local SAM, it's going to send a request up to the domain controller, and say, "hey, I have this user, he claims his name is Shad, his password is PASSWORD, what do you think?" Active Directory, the domain controller will say, yep, that's a match, and send that back, and now that I've been authenticated, I actually get a little invisible I.D. badge called an Access token. That access token can be used at any other server or computer in the domain, because they all participate with that Active Directory database. so now, So now, you know, I want to use something on B, no problem. I don't have to log in again, I've got my access token, it's good everywhere inside that domain.

Computer Roles Within a Domain
2:13-2:17
Within a domain, computers can have one of three different roles.

Client
2:18-2:29
One role is to be a client. That generally means that it's running a client operating system, like Windows 8, and it does have its own SAM. We try to avoid using that because we want to use Active Directory.

Member Server
2:30-2:50
Once you install Windows Server 2012, at that point, it's going to be what is called a Member Server. A Member Server is running a Server OS, like Server 2012, but it has its own SAM. At that point it might be in a Workgroup, or it might be a member of a domain. We have plenty of member servers within the domain that belong to the domain, but they have their own SAM.

Domain Controller
2:51-2:59
The third role might be a domain controller. Domain controllers do not have a SAM. They have a copy of Active Directory.

Domain Controller or Member Server

3:00-4:33

It's important with Windows Server 2012 to keep in mind whether your server is a domain controller or a member server. If it's a domain controller, it has a copy of Active Directory; if it's a member server, it has its own SAM. It's not too much of a problem for the domain controllers, which we will talk about in a moment, but with the member servers it is important to remember that they do have their own database. When they join the domain, they can participate in the Active Directory database, but they still maintain that SAM.

Every computer has some kind of a database, whether it's a SAM, or Active Directory. How is that relevant? Well, let me give you an example. Support; I need to create a user that's going to be used to do backups. When I go in, and I add it to the backup operator's group, on a domain controller, because the domain controllers all have a copy of Active Directory, suddenly my backup dude or whoever it is I've created, is going to be able to do backups on all of the domain controllers. However, that user account will not be able to do backups on member servers, because they have their own independent SAM. I would still need to go to each member server and add that user into the group, on the database, on that server, in order for it to get rights.

In another lesson, we'll see how to do that automatically with group policy. But, the moral of the story is, on a member server, they have their own SAM, and to get rights in that server we still have to do something with the domain accounts in order to be able to get them rights. The only great thing is, by being a member of the domain, we still can use that one account up in Active Directory.

How Domain Controllers Work

4:34-4:46

Last, let's take a look at how the domain controllers work. Active Directory is a Multi-master, Loosely-consistent database. What does that mean? This is absolutely the most succinct definition we can possibly get.

Multi-Master

4:47-5:20

Now, I'm not the world's greatest drawer, so you've got to imagine that this is the United States and I'm from the East Coast, I live roughly about there, we'll call that New York, why not. If I have a domain controller, in New York, and then we'll put one let's say in Chicago, and then we'll try to put one in L.A., all of these domain controllers can accept changes. That's your multi-master. Somebody new gets elected President of the United States no problem, you connect up to the domain controller in New York, or maybe there's one in Washington D.C., I make the, you know, President account, and any one of those domain controllers will accept changes.

Loosely-Consistent

5:21-5:55

Loosely-consistent means that at any given time, some of these domain controllers might be out of date, so when I create this President account on the domain controller in New York, these other two domain controllers are kind of behind the times, and so this domain controller that has accepted the change has to replicate the information to the other domain controllers. Information is sent across using replication. That's what makes this database loosely consistent, so at any given time, not everybody might know about everything, but replication will keep it up to speed so that generally speaking, most of these domain controllers will have a full copy of Active Directory.

Summary

5:56-6:29

As you go through and you're setting up your servers, keep in mind it's either a member server or it's a domain controller.

Domain controllers only belong to one domain at a time. If you have multiple domains, you have domain controllers in each domain, but they have a full copy of the Active Directory database for their domain. It's multi-master, so changes can be made at any domain controller in that domain. It's loosely-consistent so that when that change occurs, the domain controllers will use replication, transfer that change over to the other domain controllers, so ideally at any given point in time, all the domain controllers have an exact duplicate copy of the Active Directory database.

4.1.3 Active Directory Facts

Active Directory is a centralized database that contains user account and security information. In a workgroup, security and management take place on each computer, with each computer holding information about users and resources. With Active Directory, all computers share the same central database.

The Active Directory structure has the following components:

Component	Description
Trees and Forests	Multiple domains are grouped together in the following relationship: • A *tree* is a group of related domains that share the same contiguous DNS namespace. • A *forest* is a collection of related domain trees. The forest establishes the relationship between trees that have different DNS name spaces. Trees and forests have the following characteristics: • The *forest root domain* is the top-level domain in the top tree. It is the first domain created in the Active Directory forest. • The *tree root domain* is the highest level domain in a tree. • Each domain in the tree that is connected to the tree root domain is called a *child domain*. • A *domain tree* is a group of domains based on the same namespace. Domains in a tree: Are connected with a two-way transitive trust. Can share resources with any other domain in the forest. Share a common schema. Have common global catalogs.
Domain	A *domain* is an administratively-defined collection of network resources that share a common directory database and security policies. The domain is the basic administrative unit of an Active Directory structure. • Database information is replicated (shared or copied) within a domain. • Security settings are not shared between domains. • Each domain maintains its own set of relationships with other domains. • Domains are identified using DNS names. The common name is the domain name itself. The distinguished name includes the DNS context or additional portions of the name.

	Depending on the network structure and requirements, the entire network might be represented by a single domain with millions of objects or the network might require multiple domains.
Organizational Unit (OU)	An *organizational unit* is like a folder that subdivides and organizes network resources within a domain. An organizational unit: • Is a container object. • Can be used to logically organize network resources. • Simplifies security administration. You should know the following about OUs: • First-level OUs can be called *parents*. • Second-level OUs can be called *children*. • OUs can contain other OUs or any type of leaf object (e.g. users, computers, and printers).
Objects	Within Active Directory, each resource is identified as an *object*. Common objects include: • Users • Groups • Computers • Shared folders You should know the following about objects: • Each object contains *attributes* (i.e., information about the object such as a user's name, phone number, and email address) which are used for locating and securing resources. • The *schema* identifies the object classes (the type of objects) that exist in the tree and the attributes (properties) of the object. • Active Directory uses DNS for locating and naming objects. • Container objects hold or group other objects--either other containers or leaf objects.
Generic Containers	Like OUs, generic containers are used to organize Active Directory objects. Generic container objects: • Are created by default • Cannot be created, moved, renamed, or deleted • Have very few editable properties

Domain Controller	A *domain controller* is a server that holds a copy of the Active Directory database that can be written to. • A domain controller is a member of only one domain. • Any domain controller can make changes to the Active Directory database. • *Replication* is the process of copying changes to Active Directory between the domain controllers.
Global Catalog	The Global Catalog (GC) is a database that contains a partial replica of every object from every domain within a forest. A server that holds a copy of the Global Catalog is a *global catalog server*. The Global Catalog facilitates faster searches because different domain controllers do not have to be referenced.

The Active Directory database resides in a file called **Ntds.dit**. It is the database file in which all directory data is stored.

4.1.4 Network Model Facts

A networking model defines how network components function and interact. The three networking models used with Windows operating systems are described in the following table:

Model	Description
Stand-Alone	In the stand-alone model: • The hosts function independently. • Communication takes place using a commonly available public network such as the Internet. • The hosts are not connected by a local area network.
Workgroup	The workgroup model is based on peer-to-peer networking. In the workgroup model: • None of the hosts in a workgroup have a specific role. Hosts function as both workstation and server. Hosts in a workgroup both provide network services and consume network services. • The hosts are linked together by some type of local network connection. • Hosts in the same workgroup can access shared resources on other hosts. • Each host has a database that handles the security for the host. Drawbacks of the workgroup model are: • Lack of scalability • Lack of centralized configuration control. • Complexity of backing up data. • Lack of centralized authentication. • Lack of centrally applied security settings.
Client-Server	In the client-server model, each host has a specific role in the network. Servers provide services such as file storage, user management, and printing. Clients request services from servers. The client-server model is also known as *domain* networking in a Windows environment. Key facts are: • Domain networking uses the concept of *security principals*, which are entities such as users, computers, and resources. • A Windows domain is a collection of security principals that share an Active Directory database.

- The Active Directory database is located on one or more servers in the domain.
  - The servers running the Active Directory database are called domain controllers.
  - Hosts in a domain must run a specified version of the Windows operating system.
- Within a domain, a host can have one of the following three roles:
  - A *client* typically runs a client operation system, such as Windows 8, and has its own security database that is used when it is not connected to a domain.
  - A *member server* runs a server operating system, such as Windows Server 2012, and has its own security database that is used when it is not connected to a domain.
  - A *domain controller* runs a server operating system, such as Windows Server 2012, and has a copy of Active Directory.

In order to have grant rights on member servers, you have to use Group Policies or grant rights on the individual member servers.

Drawbacks of the client-server model are:

- Increased cost to implement due to specialized hardware and software requirements.
- Increased planning time required for implementation.

4.2 Domain Controllers

As you study this section, answer the following questions:

- To remove Active Directory from a domain controller, what action must you take before demoting the domain controller?
- What are the four methods you can use to install Active Directory Domain Services?
- How does Active Directory use the schema?
- What is the function of a Global Catalog server?
- How is a Global Catalog server updated?
- What is the function of DNS in Active Directory?
- What is the purpose of the directory partition?
- Which partition types can be included in the directory partition?

After finishing this section, you should be able to complete the following tasks:

- Create a new domain.
- Configure a Global Catalog server.
- Troubleshoot DNS issues with Active Directory.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0. Manage Active Directory.
  Configure Global Catalog Servers

This section covers the following 70-410 exam objective:

- 501. Install domain controllers.
  This objective may include but is not limited to:
  Add and remove a domain controller from a domain
  Upgrade a domain controller
  Install Active Directory Domain Services (AD DS) on a Server Core installation
  Install a domain controller from Install from Media (IFM)
  Configure a global catalog server

4.2.1 Domain Controller Installation

Domain Controller Installation
0:00-0:13
Let's talk about the process of building a domain controller. In Windows Server 2008/2008 R2 and earlier versions, you just ran a command called dcpromo, and then you walk your way through the wizard, and you'd have a domain controller.

Promotion vs Demotion
0:14-0:22
The process of making a member server into a domain controller is called promotion. If you take a domain controller and you make it back into a member server, that's demotion.

Active Directory Domain Services Role
0:23-1:19
With Windows Server 2012, it's a little bit different. First, you're going to go through and add the Active Directory Domain Services role. As soon as that's done, it's going to pop up a notification and prompt you to promote the computer to a domain controller. When you click on that link, it's going to run the Active Directory Domain Services Configuration Wizard. It's pretty much just like dcpromo used to be, but it's not dcpromo.
Once you get through the wizard, you're going to have lots of choices. You can create a replica domain controller, or another domain controller for an existing domain. You could make a new tree in an existing forest. You can make a new domain in an existing tree. You could make a whole new forest. All the choices that you could possibly have or that you had in the past you're going to be able to go through and specify. You'll also be able to opt whether you're going to install DNS. Once you get through the wizard, it'll reboot and you'll have your domain controller.

PowerShell
1:20-1:44
Now, if you don't want to go through the configuration wizard, or if you have Server Core, there's a couple other ways that you can do it. One way is through PowerShell, and that's kind of new. There's still support in Server Core for the old dcpromo command working in conjunction with an unattended install file, but it's still out there. It's really just for Server Core. The preferred method would be to use PowerShell.

Install From Media (IFM)
1:45-2:42
Now, if you're worried about replication while you're promoting the domain controller, you can actually go through and do what's called an Install From Media. You would use the ntdsutil command to create an IFM, Install From Media set. This set has everything that's in Active Directory except all the passwords. So if for some reason it becomes compromised or gets stolen even though the password's pretty secure anyway, they're still not in the IFM set. You would take the IFM set with you to the new domain controller.
When you run Advanced Configuration, you can specify the IFM as the source for Active Directory instead of sourcing it from another domain controller. Anything that's changed since you made your IFM set will get taken care of via replication. It's going to be a lot less replication than you would have if you just ran the configuration wizard and let it pull everything across the LAN or the WAN from another domain controller.

Installing Active Directory
2:43-3:25
Let's take a look at a quick recap. We're going to go into Server Manager. We're going to add our Active Directory Domain Services role and then the notification will pop up and we'll use that to run the Active Directory Domain Services Configuration Wizard. If you're much more of a command line person, you could use PowerShell.
Again, remember on Server Core, we can also use dcpromo with an unattended install file, but Microsoft recommends PowerShell as opposed to that. If we're concerned about replication, we'll create an install from media set using the ntdsutil command and then we can use that as the source while we're promoting our domain controller and let replication take care of anything that's happened since we made our IFM set.

Summary
3:26-3:40
That's the process that we're going to use to promote a member server to a domain controller in Windows Server 2012. It's pretty cool. It's not much different from what it's been before. Once we have our domain controller up and running, then we can jump into Active Directory and configure our environment.

4.2.2 Domain Controller Upgrades

Domain Controller Upgrades
0:00-0:36
We're going to talk about upgrading domain controllers. These are domain controllers in an existing environment that are running Windows Server 2008 or 2008-R2. You need to make sure that you've mastered which versions of 2008 or 2008-R2, and their service packs, and be upgraded to 2012. Remember, starting with Windows Server 2008-R2, Microsoft dropped support for 32-bit server operating systems. If you have a 32-bit 2008 server, it's not going to be upgradeable to 2012, and anything below 2008 is not going to go either.
Let's make sure we get a list on the whiteboard of the things that you need to do.

No x86 (32 bit) Upgrades from Server 2008
0:37-0:41
We'll add this one first - so no X86, 32-bit 2008 upgrades, and no upgrades from server 2003.

No Upgrades from Server 2003.
0:42-0:51
If you're not sure what can be upgraded to what, you should go and watch the Upgrading the Operating system lesson, in this course.
The next thing you've got to make sure is that your forest is at the Windows Server 2003 functional level.

Windows Server 2003 Functional Level or Higher
0:52-2:37
What that means is, you have all of your domain controllers everywhere in the forest at least running Windows Server 2003. Even though we can't upgrade 2003 to 2012, we can still have 2003 servers in the environment as we're bringing in our first Windows Server 2012 domain controller. But they've got to be at least 2003 or better, and you've got to raise the functional level. You want to make sure that you have adequate free space on the server for the upgrade. For the upgrade itself, you're going to need some free space for the operating system, but you also need to accommodate changes to the Active Directory database. The Active Directory database is in a file called ntds.dit, and it's normally in the Windows
tds folder, unless somebody's done something different. Microsoft recommends that you have at least 20% of that file in free space before you start the upgrade of the domain controller or the upgrade might fail. If for some reason you are tight for space on the domain controller, hopefully that's not the case, you could try doing an offline compact of the file in order to get space back. Very briefly, if you've ever played the game Jenga, it's a tower of blocks. You remove blocks until the tower topples. That's how databases work. When you remove things from a database, it doesn't rebuild the tower. It just leaves spaces. Normally, Active Directory will compact itself, but just in case it hasn't done that recently, or hasn't done it effectively, you could do your own offline defragmentation, and maybe get back a little space. You shouldn't be doing an upgrade of a server that's that sketchy on how much free space you have on the hard drive.
My recommendation is, you want to address that first, before you start working with the operating system. We'll go ahead and add that to our list as well.

At Least 20% of AD Database Size in Free Space
2:38-3:44
We're going to make sure that the drive that has the active directory database, ntds.dit - normally in C:\Windows t\ds, has 20% of the size of the database file in free space - before we start our upgrade. We can try to do an offline defrag to regain space if necessary, but we're really actually going to make sure we have plenty of space left so that we can upgrade it and not cut things quite so close.
When you're upgrading your first Windows Server 2012 domain controller, there are some changes that need to be made to the existing infrastructure to accommodate Windows Server 2012. Historically, what you would do is manually run a command, called adprep. Now, if you just upgrade the domain controller, or you build a new domain controller and install it as a replica domain controller, Windows Server will run adprep automatically, except for one of the commands. The idea was to try to make it easier on LAN admins, that they didn't need to run so many commands in advance. Unfortunately they left one out. We're going to look at the one command you need to run if you allow this to run automatically, and then how you could do this manually, just so you have a little bit more control over the process.

adprep /domainprep /gpprep
3:45-4:01
Adprep runs automatically except for adprep \domainprep \gpprep. The gpprep just expands and makes some additional options in the group policy objects container, and makes room for new group policy objects.

adprep /forestprep

4:02-4:32

If we're going to run this manually, there's actually three commands we have to run. The first is adprep forestprep. That has to be run on the schema master. If you're doing this automatically, then the schema master has to be available when you're promoting that first domain controller. You're going to need to use an account that is a Member of Schema Admins, Enterprise Admins, and Domain Admins. This is a pretty powerful account. The default administrator account should be a member of all of these groups, but you always want to be very careful when you're updating the schema.

Schema

4:33-5:12

The schema is the template for active directory. When you're running your forestprep, what it's doing is it might be adding classes and properties into the schema, which then every object in Active Directory will receive those new properties, so that information can be entered into Active Directory that's pertinent to Windows Server 2012.
You always want to be very careful when you're updating the schema master, because if there's a problem, it would wreck the whole forest. That being said, don't be afraid to do this. Just make sure that you run /adprep /forestprep on the server that has the schema master, and that you're a member of all three of those groups, particularly schema admins, or you're not going to have any ability to change that schema.

adprep /domainprep

5:13-5:47

Our next adprep command is going to be adprep domainprep. Adprep domainprep has got to be run on the server that has the infrastructure master. If you're doing this automatically, and you're going ahead with the /dcpromo, you need to make sure that the infrastructure master is available and can be contacted by that first server. In order to do this, you need to be a member of domain admins, and you're also going to need to run the /domainprep /gpprep to prep group policy. In a multiple domain forest, adprep /domainprep needs to be run in each of the domains. The /forestprep is run once. /domainprep is run in each and every single domain.

adprep /rodcprep

5:48-6:20

The last adprep switch is going to be run only if you're bringing in a 2012 RODC, Read-Only Domain Controller. But since it kinds of rounds out our adprep commands, I want to make sure that we're aware of it. Let's take a look at adprep /rodcprep. Adprep /rodcprep - you can run this on any computer that's joined to the domain, even a client - although that would tend to run out of server. The only key is you need to be a member of Enterprise Admins. You just run that before you bring in your first RODC, and you should be good to go.

Summary

6:21-7:00

Bringing in your first Windows Server 2012 domain controller into an existing forest is not that bad. Just make sure you've met the prerequisites. Make sure you can upgrade the operating system according to the rules for upgrading Server 2008 and 2008 R2. Make sure you've got plenty of free space on the hard drive, at least 20% of that ntds.dit, but ideally, lots more space than that. Make sure you've got your Windows Server 2003 forest functional level. Then either let it run adprep automatically and you manually run the /domainprep /gpprep, or go ahead and run the four adprep commands individually. Then you'll be able to integrate Windows Server 2012 into your environment, and you'll be very happy with that.

4.2.3 Global Catalog Server

Global Catalog Server
0:00-0:12
Let's talk about Global Catalog Servers. Domain controllers have a complete copy of the Active Directory Database or their domain, but if you have a forest situation with multiple domains, then it becomes a little tricky, so I've created a forest called northsim.com.

Log On Example
0:13-1:02
I have some trees in my forest, northsim.com being my forest root domain. I did a southsim.com tree and then here's another domain that's part of my northsim.com tree, which is west.northsim.com, and let's suppose that this user here named Shad is going to log onto this computer here, doesn't matter what the name is, but we can call it Client 1. Client 1 is going to query its domain controller in the west northsim.com domain, saying hey, I'm looking for a user named Shad. That Shad account doesn't exist in that domain, so that domain controller is going to have to contact a global catalog server.
Global catalog servers are always domain controllers, so DC1 might be a global catalog server, or it might not.

Global Catalog Server
1:03-1:24
Basically a global catalog server is used for searches and log-ons. And here's what's great about them: they have a full copy of the database for their own domain, but they also have a partial copy of all other objects in the forest. So they know about all the other objects, they don't have all the properties, but they've got some of them, so here in my log on example, DC1 in west.northsim.com is going to talk to the global catalog server and say, "hey, I need to find this user named Shad".

Continue: Log On Example
1:25-2:14
The global catalog server will have a partial copy of the Shad account, we'll know that it's in southsim.com, and have the little bit of information needed for DC1 to authenticate me and get me logged in at Client 1. That's how it helps with logons.
It also will help with searches. For example, once I get logged in here on west.northsim.com, maybe I want to go here and do something on a file server named FS1, that lives up in northsim.com. I can do a search, and my global catalog server there will know that FS1 lives in northsim.com and be able to direct me to those services.

How to Make a Global Catalog Server
2:15-2:27
To make a global catalog server, you simply go in, in Active Directory Sites and Services, and you want to go in the Properties of the Domain Controller. Anytime you're dealing with a Microsoft test, make sure you know where to go to do things, if you have things that you commonly search in the forest that are not a part of the global catalog server.

Searching
2:28-3:21
Like for example, let's say, you go in on every single user and you fill out the department, and in your company, very often people are saying "well, I need to find everybody that's in the sales department", or "I need to find everybody that's in the production department".
The department is not a field that's kept in the global catalog server. Again, global catalog servers know about all the objects in all the other domains, but just a limited subset of their properties, not very many. If I go to search by department, I'm not going to get a great list, because my domain controller will contact a global catalog server if it's not one, say "hey, I need a list of everything where the department is sales", but that global catalog server isn't going to know about any other sales department users in domains other than its own.

Adding Properties to the Global Catalog
3:22-3:56
In that case I can add that property of the user accounts to the global catalog. To do that, I've got to modify the template for Active Directory, which is the schema. I've got to go into the schema, make a change, and say hey, there's a department attribute, let's go ahead and replicate this to the global catalog, there's going to be a lot of replication initially, while all the global catalog servers get up to date. Then they will have a copy of the department, and then I can perform searches and be very confident that when I search for everybody in the sales department, I am in fact going to get back a complete list.

Global Catalog Servers in Each Site

3:57-4:31

You want to make sure that you have a global catalog server at least in each site. If bandwidth isn't a problem, there's no reason why you can't make every domain controller a global catalog server, but if replications is an issue, or you have trouble with bandwidth, you want to have one global catalog server at least in each site.

Global catalog servers are domain controllers. They have a full copy of their their own domain database, and they have a list of all the objects in all the other domains, with just a limited subset of their properties, that you use for searching, and to help people log on to get users authenticated. There should be at least one in each site.

4.2.4 Record Configuration Issues

Record Configuration Issues
0:00-0:33
Let's talk about what to do if the domain controllers aren't necessarily responding the way they need to respond. I've got just a simple domain, Northsim.com@client1, and let's say I got to log on at Client1. We know that Client1 is going to talk to the domain controller. The domain controller may have to talk to the global catalog server or may not. Once I'm authenticated, my username and password match, it's going to send me back an access token which is an invisible electronic ID badge that's going to get me in all over the domain.

Find the Domain Controller
0:34-1:22
Here's the problem right off the bat: Client1 has got to be able to find the domain controller. The way anything is found in Active Directory is using DNS. There has to be somewhere a DNS server so that Client1 can say, 'Hey, I'm looking for a domain controller on Northsim.' DNS says 'talk to DC1' and now I talk to DC1, and I get my access token, assuming I've got a correct username password combination. The way DNS knows about DC1 is, DC1 has got to contact the DNS server and register itself. Not only does it have to say, 'Hey, I'm DC1. Here's my IP address,' but it's going to register some other records that say, 'By the way, I'm a domain controller,' and we'll go into those specifics in the DNS lesson.

Ping the Domain
1:23-2:30
The deal now here we're talking about is, let's say this client talks to DNS and DNS can't refer it to a domain controller. What are we going to do? Usually, you know that there is this type of a problem because you go hit control alt delete then enter username and password and you'll get back an error message saying there are no domain controllers available to service this log on request, or something along those lines. I would say to begin troubleshooting that, you would have to log on as a local account, something that's not dependent on the domain controller, and trying pinging your DNS server to make sure it's talking.
I actually circumvent that a little bit. That's the classical answer. What I would do even before that is ping the domain. If I'm looking to get logged in to Northsim.com, I will ping Northsim.com. That accomplishes two things. Number one, it's going to connect with the DNS, and number two, since I haven't specified a particular server, it will have to go and see who is the domain controller for Northsim.com, give me that IP address, and then I'll ping the domain controller.

Missing Records
2:31-3:18
You can try to ping the domain by name, hitting Northsim.com and there's no reply. that will tell you definitively that the records that should be in DNS to locate the domain controller are missing.
The way to fix this is to go into the domain controller and restart the net logon service. Restarting the net logon service will tell the domain controller to contact DNS and re-register all the appropriate Active Directory records. Certainly, if you ping the domain and you don't get an answer, you could try pinging your DNS server by IP address just to make sure it's not a network connectivity problem. If you've eliminated a hardware problem, it's a PIP problem, and if it's a DNS problem with Active Directory specifically, you're going to go ahead and restart the net logon servers on the domain controller.

4.2.5 Creating a New Domain

Creating a New Domain
0:00-0:25
This demonstration is going to show you how to create a new domain. Before you create your domain, with your very first domain controller and because we're in a workgroup, technically right now, we are a member server, we have a SAM. The domain that I'm going to be creating will end up being the forest root domain, because it's going to bring into existence a forest, a tree, and a domain.

Computer Name
0:26-0:43
Very important to make sure that you're happy with the computer name. Once you install Active Directory, you can't change the computer name after that. Make sure that you've got a computer name that you're happy with. I've renamed mine DC1 and I went through the reboot already.

Time Zone
0:44-1:21
Make sure that your Time zone is set correctly. Active Directory is very sensitive to issues with time, in security protocol for Active directly, which I say kerberos, some people say cerberos, tomato/tomato.
It does not allow for more than a five minute clock skew between any clients in the domain. That means that a client cannot have their time be more than five minutes off from the server otherwise they won't be able to successfully talk to the domain controller.
Most common cause of that issue would be time zones, so make sure that your time zone is set.

Static IP Address
1:22-1:56
It's also recommended for a domain controller that you have a static IP address. You can see I do not have a static IPv4 address, so we're going to go ahead and set one up, because there is something important I want to show you in here for installing a new domain.
I'm going to go into the Properties. You want to use an IP address that's appropriate for your network. A Default gateway of course should be the internal address, and the router. That's not that important.

DNS
1:57-2:58
What is very important is DNS. In order to support a domain in Active Directly, you have to have internal DNS. If you don't have some kind of a DNS server that you're already using, when we install Active Directory, it's going to install DNS along with it.
The server should use itself as the preferred DNS server. I have two choices in here. One choice would be to type the exact same IP address that I've typed up for the IP address. What's even a little bit more slick is to type the loop back address, 127.0.0.1, which ensures that even if I have to come in here and change my IP address, it's just going to point to itself for DNS. It uses itself as a DNS server so that it can register all the records required for Active Directory. If it doesn't have DNS internally and it doesn't register those records, nothing is going to work.

Install Active Directory and Create a New Domain
2:59-3:04
Now, we're ready to install Active Directory and create our domain.

Adding the Active Directory Domain Services Role
3:05-3:12
I'm going to go back to Dashboard, and we need to click Add roles and features, and add in the Active Directory role.

Removing Roles
3:13-3:50
Notice, one thing that's a little bit odd with server 2012 is if I need to remove a role, I go into add roles and features and here's the link to take me into start the remove roles and features wizard, which is actually the easiest way to get into that one. Since I'm installing, I'm just going to hit Next. This is a role I'm installing, so I'm going to hit Next. I'm installing it on this server, so I'm going to hit Next. Then the role that we need to install is Active Directory Domain Services.

Adding Features
3:51-4:39
You'll notice as soon as I click that, it comes up and says, hey, you need all these features for this. Do you want to add them in; you notice it's also including the Management Tools if applicable which it is. I'm going to go ahead and click Add features because I do want all those things to get installed. With any particular role, you can also see a description of the role over on the right hand side. If you're ever confused as to which role to pick, simply click on the role, look at the description, and you'll know if that's the role you're looking for.
We're going to hit Next. It takes us to the Features. That box that popped up where I clicked add features has already selected the appropriate features that I need, so I don't need to make any changes in this list. I'm just going to hit Next. It gives me a little bit of a spiel on Active Directory.
DNS Server Required
4:40-4:58
It does say requires a DNS server. If you do not have a DNS server installed, you'll be prompted to install the DNS Server Role on this machine. It's also going to install DFS, File Replication; anything that's required by Active Directory is going to get pulled in. I'm going to hit Next.
Restart
4:59-6:03
If you want to, when you're installing a role, you can come up and check Restart the server automatically if required. If a restart is required, it will restart automatically. Do you want that? Yes.
Not all roles require a restart, but that would ensure that you don't have to wait, and be like, oh, okay. Now I notice that my role is installed. Let me restart. This will automatically restart it. I'm going to go ahead and click install.
The other thing that's interesting is that you can close the wizard and it will continue installing and give you a little flag in the notification area when it's done. We'll actually do that.
Once my feature has installed, this one didn't require a restart, so it didn't restart. Even if I didn't have this open, I would see the little notifications up here and click on it. It tells me that my installation succeeded and click on Task Details if I want. See any notifications. If I want to remove this from the list, I can hit the X. It tells me that in order to have a domain controller, I've got to Promote the server to a domain controller.
Promote the Server to a Domain Controller
6:04-6:23
This has replaced DC Promo, so there's no DC promo in the GUI.
My next step in creating a domain is to go ahead and click this link. You notice I have some choices here.
Replica Domain Controller
6:24-7:02
I can start with Adding a domain controller to an existing domain; that's called a replica domain controller. I always like to tell my students, As soon as you get into computers, you want to emulate Noah, that guy from the Bible with the boat. You want two of everything, so you want to at least two domain controllers in every domain, more is better.
We're always looking for fault tolerance, which means something went wrong, and we survive without any interruption to the end users. Having one domain controller is a problem. Because we're just doing this for demonstration, we're just going to have one. Normally, I would want at least one replica domain controller.
Deployment Configuration
7:03-7:44
If I needed to, I could add a new domain to an existing forest, and then I would be able to choose; is this going to be a Child Domain, or is it going to be a new Tree? You can see if I click Tree, it's then going to prompt me for the name of the Forest and the name of the new Tree. It's a Child domain. I would put in parent the domain name and the name of the New domain.
Since we don't have anything, we're going to have to go ahead and Add a new forest, and we will specify the forest Root domain name.

Naming Root Domain Name

7:45-9:04

Microsoft recommends that if you own a second level domain, and that's your web presence. Let's say out on the internet, my web presence is www.shad.com, and I own shad.com. They don't recommend that you name your Active Directory configuration "shad.com", because anybody who can get into your webpage is also going to know what to try to hack for Active Directory.

They used to recommend .local, but they don't recommend that anymore. What they recommend is, if I bought shad.com and that's what I'm using for my website, I would name it something like corp.shad.com or internal.shad.com, something that's a third level domain name so that hackers wouldn't know the actual name of the Active Directory. Just to make things simple for our demo environment, I'm going to name my domain "northsim.com". Of course, in real life, I'd make it "corp.northsim.com". Here, I'm just going to keep it easy so that the name stays short. Once you've decided on a name, you hit Next.

What it's doing now is making sure there's nothing else with that name.

Functional Level

9:05-10:03

Now I'm being prompted for the functional level of the new forest and root domain. If we really are creating a new forest, there's no reason not to leave both of these at Windows Server 2012.

What the functional levels do is specify the oldest operating system running on any of our domain controllers. If for some reason I anticipated that I was going to install a brand new Windows Server 2003 Domain Controller into this forest, I would need to leave my Forest functional level at 2003.

Since we're making a brand new forest, we know it's 2012, we're going to leave it 2012, and that way, the boss can't try to make us install any old junk we have laying around.

It's going to install DNS because we don't have DNS, and the very first domain controller that's creating the forest has got to be a Global Catalog server. We don't have any choices about that.

Directory Services Restore Mode (DSRM) Password

10:04-11:43

What we do have to set in here is the Directory Services Restore Mode (DSRM) password. This is the password that would be used if we had to boot the computer, and Active Directory wouldn't work. I can't get Active Directory to function. I boot the computer into the Directory Services Restore Mode, and then because Active Directory isn't running, I can't use my domain administrator password. They don't have a domain to get to the domain administrator account to give the password, so on and so forth. I would use this DSRM password, which is another hidden administrator account that's only in play if I have to boot into that mode because Active Directory isn't working. You have to give it some kind of a password.

Documentation is a wonderful thing, so you should always document any choices that you make while you're installing or performing any type of function in your environment. Now I've got my password, I'm going to hit Next.

The warning that comes up here is just saying, hey, you said this is northsim.com. I can't talk to .com to make a delegation. That's not a problem. We're not trying to integrate with the internet, so we can ignore this.

If this really were a child domain, maybe we might have to get into a delegation, and we'll learn about delegations in the lessons on DNS. When we're creating our domain, brand new forest, we can completely ignore this error and just hit Next.

NetBIOS Domain Name

11:44-12:24

Even though NetBIOS hasn't been around theoretically since Windows 2000, it still comes up with a NetBIOS equivalent name, which is going to be limited to 15 characters. As much as possible, try to keep the names of anything in your domain, your organizational units, pretty small. Sometimes, if you're using big names, either it won't work or it gets very annoying if you have to do command line or scripting. That's a fine NetBIOS equivalent name. We hit Next.

Now it tells us where it's going to store the Database, the Log files, and the SYSVOL. The database folder by default is C:\Windows\NTDS.

Database Folder

12:25-12:41

It's going to be named NTDS.dit, and the recommendation is you would want to put that database on a RAID5 array, or put it on a very fast disk. Every database uses log files.

Log Files
12:42-13:32
The purpose of the log files is to make sure that when a transaction is going through any type of database - in which Active Directory is just one flavor - that the transaction either goes through all or nothing. The log files are used. In the event of a crash, it would know how far it got with each transaction.
If possible, you could store those on a different disk, and it could even be a mirrored disk. The one most important thing is, if you're going to change the defaults, every administrator in the world is going to assume that these things are living in C:\Windows\NTDS. If you change them, make sure that's part of your documentation. You would hate to move these files to some other disks, and then somebody comes along and says, hey, we don't need the E: drive, let's get rid of that, and now you have no more Active Directory.
SYSVOL Folder
13:33-13:47
The SYSVOL folder is a share that's created during installation, and it's used to store Group Policies and logon scripts. All these things are fine, we're going to go ahead and hit Next.
Review Options
13:48-13:55
At this point, we can review the choices we've made. This is a great page, now, let me tell you why.
Installing Active Directory on a Server Core Machine
13:56-14:48
If I were going to install Active Directory on a Server Core machine, I would need an unattended install file to make that happen. Here, I can go through and click view scripts. This is going to give me a PowerShell script that I could use to do exactly the same settings that I've done in this wizard. If I did have to do this on a Server Core machine, I could go through the GUI, on a regular machine, follow this wizard through, get everything set up the way I want it, and then copy that script over to Server Core and just run it. I could save this someplace. I'm not going to save it, but that's a great way to save yourself some work with server core.
We're just going to hit Next.
Prerequisites Check
14:49-15:02
It will go through a Prerequisites Check. Some prerequisites really have to be done in order to install Active Directory, others are just optional. It's going to tell me whether or not I actually can run dcpromo.
Install
15:03-16:24
Here, I've got some errors, but they're not errors that are going to prevent me from installing it, because all the checks passed successfully and I can click Install to begin installation, and it's going to automatically reboot at the end of the promotion operation. We'll go ahead and click Install.
Now that we've installed Active Directory, I'm going to hit Ctrl, Alt, Delete. It's prompting me to logon as the administrator of this domain.
We can tell that Active Directory was installed successfully, because now I have an addition here in the list for Active Directory. I can come in and get some statistics about it. Notice it also installed DNS, and I can get some statistics about that.
Now, up under the Tools menu, I have all the different administrative tools that I need for Active Directory; Active Directory Administrative Center, Domains and Trusts, Special Module for PowerShell, Sites and Services, Users and Computers, Group Policy Management. I've created my domain successfully and I'm ready to go forward and get it set up for my environment.
That's how you create a new domain.

4.2.6 Deploying Active Directory with Windows Azure

Deploying Active Directory with Windows Azure
0:00-0:16
In this video we're going to take a look at deploying Active Directory with Windows Azure. You have some choices. You can integrate your Active Directory with Azure or just completely keep it up in the cloud. It's fairly easy to keep up in the cloud; you can just create users up there.

Creating a Hybrid Environment
0:17-0:34
What's more difficult is to maintain a hybrid environment. So that's what we're going to take a look at--having an internal Active Directory environment and having it sync up with Windows Azure, so that we can then take advantages of the resources up in the cloud.
There's a lot of things that you need to do to get this to work.

Local Active Directory Domain
0:35-1:07
First off I'm going to show you what's going on in my local Active Directory domain. In order to sync up with Azure you have to own a domain name and you have to be able to prove that. A lot of times you hear, essentially your domain name becomes a tenant of Office 365. I actually own Builditrite.com so I'm using a sub-domain corp.builditrite.com. But make sure you're aware of that you have to own the second level domain name in order to participate with the Microsoft Online Services.

Active Directory Users and Computers
1:08-1:28
Also I want to take a look in Active Directory Users and Computers before we set this up, so that you'll see the differences. So one of the things that you can see is that I do have a user named, Shadow Farrell. I actually have a group name, WorkFoldersUsers. These are some of the things that we would expect to see get copied up to Azure once we set up our synchronization.

Setting Up the Directory Synchronization Server
1:29-2:16
Now not only do you need your internal Active Directory structure, but you're actually going to need a server that will function as the directory synchronization server. So let's go over to our directory synchronization server and let me show you how we get that set up.
So we're here on a directory synchronization server. The first thing you have to do is get an account up on Office 365 that you can actually use to host Active Directory and I've done that. I've signed up for a tech named jwilson and he was given a login for this as jwilson@self632.onmicrosoft.com. I type his password. Once you get logged in to the portal you want to go into Admin and you should be able to click Office 365.

Viewing the Office 365 Dashboard
2:17-2:47
What we're looking really for is the dashboard. So the dashboard will give you information about the health, different things that you've signed up for Exchange, Link. If I go to 'users and groups,' you can see that right now my Administrator, John Wilson, is the only one that's there.

Set up Synchronization with Active Directory
2:48-7:40
So we want to go ahead and set up synchronization with Active Directory. Now you can go to setup for this and do a sort of quick start. But I'm going to run through it the way that has you do it manually. So the easiest way to get to the link is actually from 'users and groups.' If you're going to enable a single sign-on, you would set that up here. You would have to upload a certificate and we'll set up ADFS, where users can log in with their normal user accounts and just get in and get what they need.
Here, we're going to set up Active Directory synchronization; we're not going to worry about single sign-on, you don't have to have it. They have really a nice wizard that walks you through it. You can check the prerequisites here. But the biggest thing is that you need to have a PDC emulator that's running Windows Server 2003, Service Pack 2 or better.
The first thing that we need to do is specify our domain name that belongs to us and verify that it belongs to us. So step two, we go to domains and add the company's domain in. Now here it's already got self632.onmicrsoft because that's what they gave us. We can't use that there. It looks like they have added the ability to use contoso.onmicrosoft.com for testing purposes. But we can also add a domain if you own a domain. So you would add a domain and go to step one and add it in. So I'm going to add corp.builditrite.com hit Next. And because I actually have it registered with GoDaddy, it's going to make it very easy, but whatever your domain is registered with you're going to have to jump into it.

For testing purposes you can use the contoso1, but I want to kind of show you what happens, because in real life you're going to want to use your own domain. So it actually takes me out to GoDaddy where I can put in my GoDaddy information. Do a secure login. It says Office 365 is requesting to make changes. And what they're basically going to do is change the DNS entries to point certain things to Office 365. So I'm going to go ahead and accept that. And then it says great, we confirm that you own this domain, you can go forward and I click finish.
Now that I've specified the domain name and confirmed ownership step two asks me to add users and licenses. But I actually don't need to do that. So I can say I'm going to add users later and hit next. Because we're going to use Active Directory synchronization. But I do want to go ahead and create the DNS records. So I'm going to move to step three. I can have it create DNS records for whatever I want to use Office 365 for. I'm just going to say Exchange and hit next and if it's GoDaddy it will just set up the records and these are the records that it actually creates. So the nice thing about once you've claimed the domain is that Office 365 can set up any records that you need. If I went and looked in my GoDaddy account, I'd see all those records having been created there. So I click Finish and you can see the domain that I own has been added to this.
Now I'll go back to 'users and groups.' Click on Set up Active Directory synchronization and continue through these steps. So now I want to activate Active Directory synchronization so I click on Activate. It says are your sure? Yes. Now after you activate synchronization the synchronized object may be only edited on premises. Which means basically I'm going to make all my changes in Active Directory at my company and then synchronize and push those changes up to the Office 365 server. I'm going to hit Activate. And now, I need to install and configure the Directory Sync tool on a Directory Sync server.
So I'm going to download that. And now that the tool is downloaded we want to go ahead and install that software. I'm going to minimize Server Manager. You can see that it requires the .Net framework 3.5 and 4.0, so we need to add those features to our server. I'm going to go back do dashboard, Add roles and features. And add in the .Net framework 3.5. 4.0 is already installed by default. It'll accept the 4.5 even though it wants 4.0. So this is the only one I have to add in. It actually doesn't reside on the server, so if you want to specify an alternate source you can do that. I'm just going to click Install and it's going to download it from Microsoft.

Windows Azure Active Directory Sync

7:41-10:04
Once the installation has succeeded I'm going to hit Close. And now we'll go ahead and run that .exe again. You can see now it runs successfully. It says welcome to the wizard. I say Next. You've got to accept the software license agreement. I can accept the default installation folder. And now you can see it's setting up. Once the installation is complete I can hit Next and then it's going to finish the wizard and dump me into the configuration wizard. So I'll hit Finish and that Wizard should open up and I hit next.
Now I have to put in my Azure Active Directory administrative credentials. This is going to be the password that can go through and get up into Office 365 so that's going to be jwilson@self632.onmicrosoft.com and then I hit Next.
Now I have to put in my credentials that have rights on my organization's Active Directory. So that's going to be my corp\administrator.
And when you do this, you need to make sure that your Active Directory is named exactly like the tenant you specify in the cloud. So if I'm corp.builditrite.com I have to get that domain up on Office 365. If I can't get builditrite.com. I can't get anything else; those names have to match.
So I type in my credentials. Then it says, "Some features require data be written from Azure back to your on-premises." We need to grant right access to the directory sync tool and I can click yes, which really allows the data to go both ways between Azure and Active Directory on site and vice versa. Now here I can also enable Password Sync, which allows the user passwords from Active Directory to sync up to Azure, so that they don't have to have different passwords out there. So I'm going to go ahead and Enable Password Sync. That's complete. Then when I click Finish it'll perform the initial synchronization.
Now I can go up and verify this by going into my Office 365 account.

Verify Changes

10:05-11:45
You can see that when I click on 'users and groups' I can see that that user from my Active Directory has been pulled up to the cloud. And I can go ahead and select that user.
Right now it doesn't have any licenses associated with it, but I would see it up there. I would then need to grant licenses to the user for whatever it is they're going to be using from Office 365, whether it's Directory Rights Management, Office Pro, Exchange. I can also see some of the details in there that it's pulled in from Active Directory, specifically just my UPN and my display name. Up under settings, I can assign the user a role and give them some rights up in Azure, but I'm not going to. They are allowed to sign-in and access services. Right now, because I didn't give them Exchange, they don't have an Exchange e-mail address. And then I can add some additional self-service options.
I'm not going to save the changes. Sometimes you have to go ahead and activate the users, so that they can be used up there, but it looks like it's copied it up into the cloud. And if I go through I should be able to-- I would see all

my users up there. I can see my domains, whether or not it's active. If I want to create a user that just lives up in Azure I can hit plus (+) and I can create a new user.

Create New User Accounts in Azure

11:46-12:21

But they'll only be up there, they wouldn't actually have an Active Directory account in my server at home setup. I can also do a bulk import here to add a bunch of users all at once. If I want to see security groups I can come up here to security groups and you can see that the WorkFoldersUsers group has been synchronized up with Azure. So once you build that directory sync server it's going to push Active Directory objects up to the cloud and I can use them for things up in the cloud.

Recommendation

12:22-12:31

I really should be making changes in Active Directory and then re-synchronizing the information on Active Directory, pushing it up through the DirSync folder.

Summary

12:32-12:43

But that's essentially the steps you would go through to allow synchronization between Active Directory on-premises and Active Directory up in Azure. And again if you don't want that synchronization you could just have users that live up in the cloud.

4.2.7 Domain Controller Installation Facts

Promoting a domain controller refers to the process of installing Active Directory on a member server. The following list contains the requirements for installing Active Directory Domain Services (AD DS):

- You must have membership in the Domain Admins, Schema Admins, and Enterprise Admins group.
- You must have properly configured static IP addresses and Domain Name System (DNS) server addresses.
- You must verify that a DNS infrastructure is in place on your network before you add AD DS to create a domain or forest.
- Use local, fixed disks for the volumes that store the database, log files, and SYSVOL folder for AD DS.
- For added security, place the database and log files on a volume with the NTFS file system.
- Because Active Directory is time sensitive, ensure that the time zone and time are correct.

After installing Active Directory, you cannot change the name of the server.

There are four methods for Active Directory Domain Services (AD DS) installation:

Method	Description

Active Directory Domain Services Installation Wizard	AD DS installation using wizards requires the following actions:
	• For Windows Server 2008 and 2008 R2:
	  ▪ In Server Manager, run the Add Roles Wizard to install the Active Directories binaries.
	  ▪ Run **dcpromo.exe** when prompted. The **dcpromo** command launches the Active Directory Domain Services Installation Wizard. This wizard can be used to install new 2008 forests, trees, domains, and domain controllers.
	• For Windows Server 2012:
	  ▪ In Server Manager, run the Add Roles and Features Wizard to install the Active Directory Domain Services role to the server.
	  ▪ You will then have the option to **Promote this server to a domain controller**. Select this option to launch the Active Directory Domain Services Installation Wizard. This wizard can be used to install a Windows Server 2012 domain controller on a new or existing forest, tree, domain, or domain controller.
	  ▪ Server Manager begins every domain controller promotion with the Deployment Configuration page. The deployment operation you select determines the options and required fields you see during the rest of the wizard.
	To remove Active Directory, you can open Server Manager, click the Manage menu, and then select **Remove Roles and Features** and uncheck the Active Directory Domain

	Services role. You then have the option to demote the domain controller.
Command Line	In Windows Server 2008, use the **dcpromo** command combined with unattended installation switches and parameter values to create forests, domains, and domain controllers. For a complete list of unattended installation switches—including default values, allowed values, and descriptions—type **dcpromo /?:Promotion** at the command prompt. In Windows Server 2012, you can install roles and features using the **Install-WindowsFeature** cmdlet. For example, use the following to install the Active Directory Domain Services role, including management tools: **Install-WindowsFeature -Name AD-Domain-Services -includemanagementtools** In Windows Server 2012, the **dcpromo** unattended operation is replaced by the ADDSDeployment module for Windows PowerShell. When using Windows PowerShell to promote a domain controller, you can use one of the following three cmdlets: • **Install-AddsForest** • **Install-AddsDomainController** • **Install-AddsDomain** When you run any one of these cmdlets without parameters, you will be prompted for the values needed to set up the domain controller. To easily obtain a PowerShell script, install a domain controller on Windows Server 2012 with a GUI and then export the PowerShell script.
Answer file	An *answer file*, also referred to as *an unattended install file*, is a list of Active Directory configuration values in a text file which is used to install AD DS on either a full installation of Windows Server 2008 or a Server Core installation. To create an answer file, you can: • Run the Active Directory Domain Services Installation Wizard and export your choices to a file. • Create or edit the answer file directly using a text editor. To perform the install using the answer file, run **dcpromo /unattend:C:\unattend.txt**, using the name of the answer file you created. Using **dcpromo** with an answer file will work in Windows Server 2012; however, it is not recommended and a warning message will be generated.
AD DS installation from media	Install from media (IFM) is an alternate method of AD DS installation. The media contains the unattended installation parameters which will create additional domain controllers, as well as the Active Directory database.

	During installation, the Active Directory database is copied from the media instead of replicated from another domain controller. Use the media installation method if you need to perform a domain controller install where the domain controller will not be able to contact another domain controller during installation. Use one of the following to create the installation media: • Run **ntdsutil.exe**. • Run Windows Server backup in Windows Server 2008 or Windows Server 2012. A critical-volumes backup includes all files on the volumes that are required to recover AD DS, which requires significantly more space than is required for AD DS installation. To install a domain controller using media, use one of the following methods: • In the Active Directory Domain Services Installation Wizard, use the **Install from Media** page to refer to the location of the shared folder or removable media. • Use the **/ReplicationSourcePath** parameter during an unattended installation to specify the location of the shared folder or removable media.

When upgrading domain controllers to Windows Server 2012 or Windows Server 2012 R2, know that:

- Only Windows Server 2008 and 2008 R2 domain controllers can be upgraded to Windows Server 2012 or Windows Server 2012 R2. Upgrade from Windows Server 2003 is not supported.
- The domain controllers must have the latest service packs installed and the editions must be the same.
- Windows Server 2012 and Windows Server 2012 R2 do not support 32-bit operating systems; this means that you cannot upgrade a 32-bit domain controller.

Active Directory on Windows Azure

Active Directory domain controllers are typically deployed on physical hardware in a data center. However, they can also be implemented in the cloud using Windows Azure. Using a cloud-based Active Directory deployment on Windows Azure would be beneficial for:

- Improving authentication performance at remote locations where a WAN link is not a suitable option and the cost and lack of technical expertise do not allow for an on-site domain controller.
- Providing a disaster recovery site.
- Deploying network applications.

Deploying Active Directory via Windows Azure can be done in two ways:

- Implementing Active Directory domain controllers on Windows Azure virtual machines (VMs) in the cloud
- Using the Windows Azure Active Directory SaaS cloud service

Using the first option, you deploy virtual machines in the Azure cloud, install Windows Server on those virtual machines, and then make them domain controllers. These cloud-based domain controllers are connected to an organization's local data center. Below is an example diagram of this option:

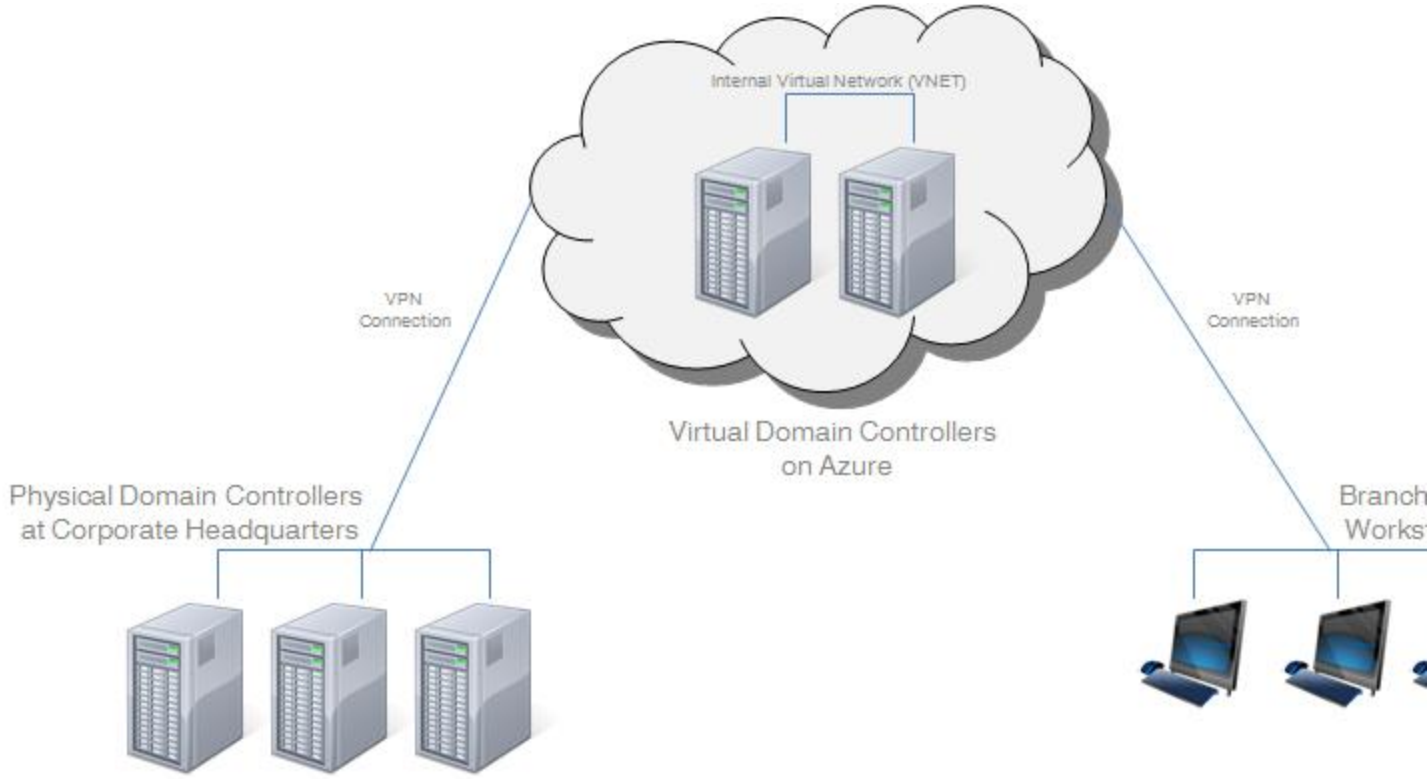

In the figure above, two cloud-based domain controllers are running on Azure virtual machines. These VMs are connected using the Azure virtual network (VNET) that is connected to an organization's local datacenter using a virtual private network (VPN) connection. The Azure domain controller VMs function as if they were local domain controllers located on a different subnet somewhere else in the network. In this scenario, users at a branch office without a locally-installed domain controller can authenticate using the cloud-based domain controllers over the Internet. No expensive WAN connectivity to the home office is required.

When using cloud-based virtual domain controllers, you should be aware of the following:

- User authentication requests should go to the cloud only when necessary. This is because an Internet connection is slower than the local network connection. For example, in the figure above, users at the company's home office should always use the local domain controllers for authentication unless they are unreachable for some reason.
- Cloud-based domain controllers should be installed in their own Active Directory site. This allows you to configure how often replication occurs.
- With Windows Azure, you are charged for outbound traffic but not for inbound traffic. To reduce outbound traffic, consider deploying cloud domain controllers as Read Only Domain Controllers (RODC).
- Running domain controllers on Azure virtual machines requires that you manually configure your own cloud-based DNS servers.

Alternatively, you can also use Azure Active Directory. Because Azure Active Directory is provided as a cloud service, there is no need to set up virtual machines and install Windows Server. An example deployment is shown in the diagram below:

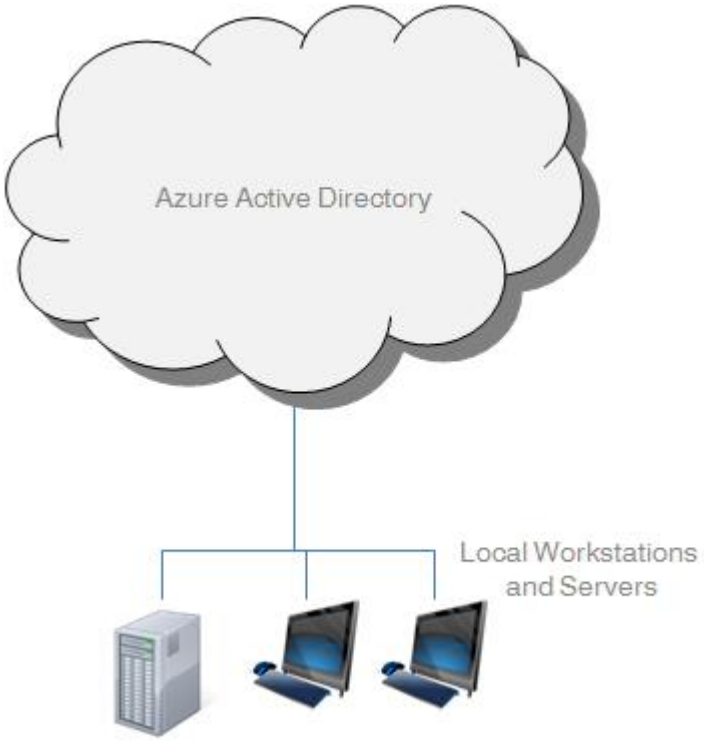

This deployment option is particularly useful because it can be leveraged to give users single sign-on access to other SaaS applications, such as Office365 and SharePoint. You can use the cloud directory as your organization's sole directory service or you can connect your local Active Directory implementation to the cloud directory service. However, be aware that the cloud directory uses a simplified schema and doesn't support Group Policy. Therefore, it is recommended that you use Azure Active Directory in conjunction with a locally-installed Active Directory deployment.

4.2.8 Domain Controller Facts

A domain controller stores the Activity Directory database for the domain in which it is located. The domain controller responds to authentication requests as well as performing other Active Directory functions for the domain.

Domain controllers share information with domain controllers in other domains; this is referred to as *replication*. The *directory partition* is used to replicate domain information. Each domain in the forest has a separate partition in the directory partition. The following table describes partitions within the directory partition:

Partition	Description
Domain partition	A *domain partition* stores the user, computer, group, and object data for a domain, as well as the domain's schema and configuration data. The domain directory partition is replicated only to another domain controller in the same domain and to global catalog servers.
Schema partition	The Active Directory *schema partition*, referred to as the *schema*, contains a definition of each object class and the attributes of the object class that can exist in an Active Directory forest. Active Directory uses the definitions in the schema to store, retrieve, and replicate data. A schema partition for a domain is replicated to all domain controllers in the forest.
Configuration partition	The Active Directory *configuration partition* stores configuration objects for each domain in the forest. A schema partition for a domain is replicated to all domain controllers in the forest.
Application directory partition	An *application directory partition* contains application-specific data created by applications and services. An application directory partition: • Is typically created by the application that will use it. • Can contain any type of object except for a security principal. • Is replicated only to specified domain controllers. • Provides fault tolerance for the application directory partition through replication.

When working with domains, be aware that:

- A domain controller can store one or more application directory partitions.
- Domain controllers can be global catalog servers and operations masters.
- If you own a second level domain and you use that as your web presence, Microsoft recommends that you name your Active Directory configuration using a third-level domain name that is reflective of

your second level domain name, but different in some respect. For example, if your web presence is **www.bikes.com**, an appropriate third-level name might be **corp.bikes.com**.

The following table describes additional functions or roles that domain controllers can have.

Function/Role	Description
Replica domain controller	Adding a domain controller in an existing domain creates a replica domain. A replica domain provides fault tolerance in the event that the domain controller fails. • Having at least one replica domain controller is best practice. • To create a replica domain controller, choose **Additional Options** in the Active Directory Domain Services Configuration Wizard and select the server you want to replicate from.
Global Catalog	The Global Catalog (GC) server is a domain controller that contains a partial replica of every object from every domain within a forest. A Global Catalog server: • Facilitates faster searches and logon. • Has a full copy of its own domain. • Has a partial copy of the attributes of other objects in the forest. • Is built and updated automatically by the AD DS replication system • Can be edited to add or remove attributes from the schema to facilitate searches. Use Active Directory Users and Computers or Active Directory Sites and Services to designate a global catalog server.
Operations Master Roles	Operations master roles, also referred to as Flexible Single-Master Operation (FSMO) roles, are specialized domain controller tasks assigned to a domain controller in the domain or forest. Operations master roles are useful because certain domain and enterprise-wide operations are not well suited for the multi-master replication performed by Active Directory. A domain controller that performs an operations master role is known as an *operations master* or *operations master role owner*. The following roles are forest roles, meaning that one domain controller within the entire forest holds the role: • The *schema master* maintains the Active Directory schema for the forest. • The *domain naming master* adds new domains to and removes existing domains from the forest.

	The following roles are domain roles, meaning that one domain controller in each domain holds the role: • The *RID master* allocates pools or blocks of numbers (called relative IDs or RIDs) that are used by the domain controller when creating new *security principles* (such as user, group, or computer accounts). • The *PDC emulator* acts like a Windows NT 4.0 Primary Domain Controller (PDC) and performs other tasks normally associated with NT domain controllers. • The *infrastructure master* is responsible for updating changes made to objects. As you install or remove domain controllers, you will need to be aware of which domain controllers hold these roles.

4.3 Sites

As you study this section, answer the following questions:

- How does a site differ from a domain?
- What is the purpose of a site link?
- What does the term "well-connected" mean when referring to networks?
- How are sites used in Active Directory?
- How do IP address and subnets relate to sites?
- How are dynamic site assignments made?

After finishing this section, you should be able to complete the following task:

- Create and configure a site.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0. Manage Active Directory.

This section covers the following 70-410 exam objective:

- 500. Install and Administer Active Directory.

4.3.1 Site Overview

Site Overview
0:00-0:43
We're going to have just a brief overview of sites. Sites aren't really part of the objectives, but we didn't want to miss anything big like sites. When we're setting up our forest, our trees, our domains, we're really looking at the logical structure of active directory, and that's going to mirror the logical structure of our organization, how it's divided up into companies or sub-companies and that type of thing. Later on in a different lesson we'll get into organizational units where we can actually go through and specify departments, and so on and so forth.
The way we talk about the physical structure of our environment is done through sites, so a good rule of thumb is, every time you cross a WAN link, you have a different site.

Sites Direct Clients to Local Resources
0:44-1:21
Sites direct clients to local resources, so for example, let's just pretend this is the U.S. I'm sitting in New York, there's a domain controller here, and there happens to be a domain controller over here in L.A. There's no reason for Active Directory to tell me, "Hey, go ahead and authenticate using the domain controller in L.A.," when there's a domain controller right in my location. By setting up sites, it's going to direct me, the client, to local resources such as domain controllers. Those are the big ones, but there are other site-aware applications like DFS or Exchange.

Sites Control Replication
1:22-1:52
The other thing that sites do are control replication, so maybe in New York I've got five domain controllers, in L.A. I've got five domain controllers. We can let these five replicate quite a bit. We don't want them crossing the WAN link, typically not at 9:00 a.m. Eastern time when everybody's logging in or at 9:00 a.m. Pacific time when everybody's logging in, so we can control replication between the sites so that we're using our WAN links really effectively.

Sites Mirror Structure of the Network
1:53-1:56
Again, the sites are going to mirror the physical structure of the network.

IP Address
1:57-2:45
Now, you might be wondering, how does the computer know which site I'm in? Well, if you think about it, the way you know which site you're in is usually based on IP address, because each of those physical locations within your company are going to be assigned a different network ID. You can go into Sites and Services and put in the different network IDs in use in your company and associate them with the correct site. Maybe if I'm in New York, that might be Network 192.168.10. If I'm in L.A., maybe that's Network 192.168.20. The computer knows if my IP address starts with 192.168.1, I'm in New York, if it starts with 192.168.2, then I'm in L.A., then it will direct me to resources accordingly.

Summary
2:46-3:12
You should be aware of sites. They need to be set up. They're probably some of the earliest things that you should set up, and they're going to become important for us when we get into group policy. Just keep in mind that sites mirror the physical organization of our network: they're used to make sure that clients get local resources and to control replication, and if you've got a feel for that, then you'll be in great shape when we get to group policy.

4.3.2 Configuring Sites

Configuring Sites
0:00-0:03
In this demonstration, we're going to take a look at how to configure sites.
Active Directory Sites and Services
0:04-0:23
Anything that happens at a site level should be done in Active Directory Sites and Services. The easiest way to get into it is through this Tools menu. If for some reason you don't like the Tools menu, once you've installed Active Directory, it will also be on the Start menu, so I can click Sites and Services that way.
Site Arrangement
0:24-1:30
Sites should be set up to mirror the physical arrangement of the company. Since the computer doesn't know how your company is physically arranged, it just creates one site to begin with that's named Default-First-Site-Name. You should go in and change this to whatever the name of your first site is.
Let's say, for example, my first site is in New York. I might change it to New York. By default, my first domain controller is going to be inside the server's folder in that initial site. That's fine if that's the way it's set up. I would go through and create additional sites if I have additional physical locations.
Let's say, for example, my company also has a site in LA. If I have multiple sites, they have to be connected by a WAN link. These site links are set up to mirror our WAN link. By default, there's one called DEFAULTSITELINK. Again, I would change the name of that to reflect New York, LA. But it's going to tell Active Directory how the sites are physically connected. Now it's prompting me. Make sure that it's linked as appropriate.
Subnets
1:31-1:43
The biggest thing is to add subnets. Once I've got my sites, the way the computer knows which site a client is in is by its IP Address. So, I need to set up subnets.
Purposes of Sites
1:44-1:45
The purposes of sites are really two-fold.
Users are Always Directed to Local Recourses
1:46-1:56
One is to make sure the users are always directed to local resources. If I'm at a Windows 8 client that's booting up in New York, I should be sent to DC1, not some server in LA.
Control Replication
1:57-2:17
The other thing that sites do is control replication. If I don't want replication to happen because it's a business time; it's 9 AM in them morning, I can set that up in here in this site.
The biggest things is to create my sites, make sure my servers are in the appropriate site, and then create a subnet that says which IP addresses are associated with which sites.
New Subnet
2:18-3:20
We're going to go ahead and make a New Subnet. Notice it could be either IPv4 or IPv6. Again, anytime users are not being directed to local resources, your problem is that you have not set up the subnet in Sites and Services whether it's IPv4, IPv6. When we put in the subnet, we need to put in the network ID, not any individual IP address. If it's IP Version 6, instead of a subnet mask, they use the term Prefix.
Either way, if it's IPv6, it's Prefix; if it's IPv4, we just put the network ID. I'm going to put in 192.168.40.0 using a 24-bit subnet mask so that anything in the 192.168.40.0 network is associated with New York.
That's how easy it is. You put in your subnets, and now anytime a computer boots up in New York it's going to have a 192.168.40.something address. Active Directory will see that address and say, "oh you're in the New York site. Let me direct you to servers and resources that are in that site".
Inter-Site Transports
3:21-3:36
If I did need to change the site links, it's done in inter-site transports. You can simply go through in the site link and say which sites are linked by that site link. Here's where I can configure my replication when it's allowed to replicate. That's how you set up sites.

4.3.3 Site Facts

Active Directory uses forests, trees, and domains to represent the logical organization of the network. Sites and subnets represent the physical layout of the network:

- A *site* represents a group of well-connected networks (networks that are connected with high-speed links).
- A *subnet* represents a physical network segment. Each subnet possesses its own unique network address space.

You should know the following about sites and subnets:

- Sites should be set up to mirror the physical layout of the company.
- Sites direct clients to local resources.
- Sites are used to manage Active Directory updates, referred to as *replication*, between locations.
  - All Active Directory sites contain servers and site links.
  - *Site links* specify how the sites are physically connected.
  - Sites allow for the efficient resource use during replication.
- You will use Active Directory Sites and Services to identify network IDs and their associated sites.
- Site assignment is made dynamically, according to Internet Protocol (IP) address and subnet mask.

4.4 Organizational Units

As you study this section, answer the following questions:

- What objects can an organizational unit contain?
- How is an organizational unit different from a generic container?
- What are the advantages of placing computer accounts in organizational units rather than the Computer container?
- How does inheritance affect child organizational units?
- How does object-based delegation differ from task-based delegation?
- How can you protect objects from accidental deletion?

After finishing this section, you should be able to complete the following tasks:

- Create organizational units.
- Create nested organizational units.
- Prevent accidental deletion of OUs.
- Delegate administrative control.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Manage Active Directory.
  - Create Organizational Units (OUs)
  - Delegate Administrative Control

This section covers the following 70-410 exam objectives:

- 102 Configure servers.
  This objective may include but is not limited to:
  - Delegate administration
- 503 Create and manage Active Directory groups and organizational units (OUs).
  This objective may include but is not limited to:
  - Delegate the creation and management of Active Directory objects
  - Manage default Active Directory containers
  - Create, copy, configure, and delete groups and OUs

4.4.1 Default Containers and Organizational Units

Default Containers and Organizational Units
0:00-0:03
Let's talk about containers and organizational units.
Containers
0:04-0:43
Containers are objects that come with Active Directory. They look a little bit different then organizational units. What gets a little confusing with the vocabulary is, both containers and organizational units can contain other objects, but containers are created by Active Directory. There's a bunch of containers that come when you install Active Directory. There is only one organizational unit that is created when you install Active Directory, which is the domain controllers organizational unit. Aside from the fact that container objects are created by Active Directory, they also differ from organizational units in a couple of other ways.
Can't Rename or Delete
0:44-0:54
One way is, you can't rename them or delete them. Really, the old adage, "if it isn't yours, don't touch it". They're really more to be used by the operating system then to be used by you.
Group Policy Can't Be Applied to Them
0:55-0:57
They also cannot have group policy applied to them.
Organizational Units
0:58-2:03
When we go in and start building our organizational structure, what we want to do is make our own organizational units. That's how we're going to set up Active Directory for administration.
Let's say we have a domain -- Northsim.com. It's going to have some container objects in there already. For example, there'll be a container in there called computers. There'll be a container in there called users, that sort of thing. There'll also be one OU already in there called Domain Controllers, and any member server that you promote to a domain controller, its computer account will be moved into that organizational unit. Now, with organizational units that you create yourself, you're going to create them to mirror the structure of your organization. Let's say I happen to have a sales department, I have a production department, and I have a shipping department. I want to mirror the logical structure of my organization. The two overriding decisions for making OUs would be, first, how you plan to administer Active Directory.
Administrative Rights
2:04-3:13
Before Active Directory came in, we didn't have any way to divide up the administrative pie, so to speak. If I was a domain admin, I had rights all over the domain, and the person who worked for the company five minutes had as much power as somebody who had worked for the company for five years. There was no way to delegate out responsibility and give people administrative rights in pieces of the network. Organizational units serve that function, so I want to develop a structure of organizational units, that's going to make it easy for me to delegate out administration of those units however I plan to support them (technically speaking).
If I plan to have a team of LAN admins that's going to support sales, a team of LAN admins that's going to support production, maybe I'm going to make OUs based on that. If I'm going to do my administration differently, you know, maybe I'm going to go by buildings or something like that, then I can certainly make organizational units that map to that, but I'm not going to let the physical layout of my network drive my organizational unit structure. It's more how I plan to support these people, because later on in a different lesson, we'll talk about delegation of control, where we can take an administrator or a group of administrators and say, hey from this organizational unit on down, you are the LAN guy.
Delegation of Control
3:14-3:41
Anywhere else in the domain you do not have any rights.
The first principle for designing our OU structure is how we plan to support the users and the computers, keeping in mind that we're going to be giving administrators control over particular functions at particular OU levels.
Group Policy Objects (GPO)
3:42-3:53
The second thing that we're going to keep in mind with our OUs is group policy- group policy objects. Our two main themes are centralized security and centralized administration.

Centralized Security
3:54-3:59
Centralized security is where we're going to put our user and computer accounts. That's what's going to drive our security.
Centralized Administration
4:00-4:50
Administration, we're going to attach group policy objects to our OUs that will make changes to the computers or the user environment based on where those computers and users are located in my structure. I'm going to design my organizational units like that.
Because container objects like computers and users can't be deleted and they have these nice names, I mean computers says computers, why don't I put my computers there? The problem with that is I can't apply a group policy object to the computers container which means I can't customize the settings on the workstation from a central location. I then have to visit the workstation and make those changes myself and that's not working efficiently. I definitely don't want to put my computer objects in the computers container. I don't want to put my user objections in the users container. Instead, I'm going to create organizational units for those objects.
Protect From Accidentally Deletion
4:51-5:23
When you're creating an organizational unit, as we're going to see in the demo, by default, there's going to be a checkmark that says "Protect from accidental deletion". The idea being that, once you have your OU structure set up, if somebody who didn't understand it were to come in and accidentally delete an OU, say you have an OU named Desktops, all your desktop accounts are in there, somebody deletes that, it would be pretty catastrophic. By default, there's going to be a checkmark that says let's protect this from accidental deletion and you'll have to jump through some hoops in the operating system in order to turn that off.
Distinguished Names
5:24-6:03
The last thing we need to talk about, in terms of OU and containers, are distinguished names. When you're working with objects in Active Directory from a command line or a PowerShell, you have to know how to refer to that object by name and by its full name. Otherwise, you won't be able to execute the commands correctly.
Let's take a look at an example of some distinguished names. Here we are with our Northsim.com domain, but I've made a few changes to the structure so that we have a little bit more to play with. I've made an OU named Sales, so this is an OU. I've made another OU inside that OU named Desktops. You certainly can, when you make an OU inside of another OU, we call that nesting an OU.
Nesting an OU
6:04-8:12
You certainly can nest OUs. The recommendation is not to go too deep, three to five is good. If you have more than five inside of each other, you probably have a design flaw.
In here I have a computer. This is an actual computer named Client 1. For comparison's sake, I've also used the Users container. Remember this is not an OU, this is a container. I did make a user named Shad in there so that we can look at the difference with a distinguished name. A distinguished name is the full name of the object, from the object all the way back to the end. If I were doing the distinguished name of Client 1, the common name is Client 1. It is in an organizational unit named desktops; you put commas between each section. That is in an organizational unit named sales. I'm going to go to the next line, but if you were doing this in a command, it would be all one line. That's in a domain container named Northsim in another domain container named com.
However, many names separated by periods you have in your domain, you have to do a DC for each of those. If this were West.Northsim.com, it would be "DC = West, DC = NorthSim, DC = com". This is the full distinguished name of this object. If there are any spaces, which you should try to avoid in naming things, then you would put the whole distinguished name in quotes so the computer knows not to pay attention to the spaces. If you do make something in a container object, which I don't recommend, but certainly you can do it, it's done a little bit differently.
Here, the common name of my object is Shad -- that hasn't changed -- but since Users isn't an OU, I'm not going to use OU. I use CN which now means container name = users in a domain container named NorthSim in a domain container named com. The only difference here is if it's in a container the type is CN for container name. If it's in an OU, it's OU for OU.

Summary

8:13-8:54

Active Directory, when it installs, is going to create some container objects. These come from the operating system. They can't be deleted. They can't be renamed. They can't take group policy. You can use them, but you shouldn't. You should create your own organizational unit structure. The only OU that comes with Active Directory is an OU named Domain Controllers, and when I promote a member server, that computer account is going to move into the domain controllers OU.

I'm going to create OUs based on the logical structure of my company, keeping in mind how I plan to support the users and computers, and how I plan to design group policy. Then, if I need to, I can write out the full distinguished name. That's really the meat and potatoes of designing an Active Directory organizational structure.

4.4.2 Creating Organizational Units

Creating Organizational Units
0:00-0:04
In this demonstration, we're going to take a look at how to create Organizational Units.

Active Directory Users and Computers
0:05-0:20
Anything that happens at a domain level is done in Active Directory Users and Computers. When I go in to create my Organizational Units, that's the snap in that I'm going to use. I can get it off the Tools menu or I could get it off the Start menu, whichever is more convenient for me.

Inside the Domain Folder
0:21-1:02
Inside my domain, I'm going to see the Default Containers and Organizational Unit that comes with Active Directory. I haven't done anything in here at all. You'll notice that the folders that had plain folders, these are containers. They cannot be deleted; they can't be renamed; they cannot have group policy applied to them.
Anything that has an icon like this, little something extra in the folder, that's an Organizational Unit. The only Organizational Unit that comes with Active Directory is the Domain Controller's Organizational Unit. Whenever I promote a member server into a Domain Controller for this domain, the computer account for that member server will be moved into or created inside a Domain Controller's OU.

Using Organizational Units Instead of Containers
1:03-1:32
Generally speaking, you shouldn't use the containers. You simply join a computer to the domain. It will end up in the Computer's container. Since there's no group policy applied to that, it's not very convenient for us. The User's container as the default Administrator account in a lot of groups. You should not use this container to make users. What we need to do is come in and create our own Organizational Units that reflect how we plan to support our users and how we plan to design group policy.

Creating an Organizational Unit
1:33-1:41
Two different ways to create an Organizational Unit; if you're a right clicker, you can right click and make a New Organizational Unit. We simply give it a name.

Protect Container from Accidental Deletion
1:42-2:23
Notice by default, there's a check mark in Protect container from accidental deletion. The container objects can't be deleted. OUs can. What we don't want is, if we make a sales OU and then someone comes in and they're new to the company, they say, "Oh I don't think we do sales anymore". They delete it, and now everything in that Organizational Unit is lost.
So by default, this is checked. That's fine. It just means that if you need to delete an OU, you're going to have to jump through some extra steps. We'll go ahead and create a Sales OU. Now I have an Organizational Unit. If you prefer the Tool bar, this button up here creates a new Organizational Unit. I can create the Organizational Unit in my domain.

Nested Organizational Unit
2:24-3:10
I can also have nested Organizational Units; which means an Organizational Unit inside of another Organizational Unit.
If I'm going to create a container for desktops, I have to make a decision. Is it going to be all desktops in my domain, or am I going to have an Organizational Unit that's just for the desktops in the sales department? If I wanted to do that, I could create a nested OU. You don't want to have too many OUs nested inside of each other. Microsoft recommends somewhere between three to five, depending on which book you read.
I think that's fine. Every once in a while I meet an author who says ten; ten is too many in my opinion. If you've got to go deeper than five, I think you might want to stop and say, is this really the best way to design my Active Directory organization? That's how you create an Organizational Unit.

How to Delete an Organizational Unit

3:11-4:17

What I do want to also show you is if you had to delete it. Because we left the Protect from accidental deletion checked, if I right click this and hit Delete, I can also hit Delete on the keyboard. Are you sure you want to delete it? Yes. It comes to me and says, "You do not have sufficient privileges to delete desktops, or, this object is protected from accidental deletion.

Unfortunately, most people don't read dialog boxes. They just see "You don't have sufficient privileges". And they say, But wait, I'm the Domain Administrator. I have as many privileges as I could possibly have. If I right-click, and I go in here to Properties, and I look at all my tabs, there's nothing in here about Protect from accidental deletion. That's what makes it a little tricky.

If you leave that checked and you do need to delete something, you first have to go up under the View menu and choose Advanced Features. That's going to show me a lot more container objects. The reason these are not shown by default, is Microsoft doesn't want you in here. You don't need to be in LostAndFound ... ForeignSecurityPrincipals ... these are used by the operating system. In very rare circumstances you might have to come in here, but on a day to day basis, you really don't need to see them.

Advanced View

4:18-4:52

Once you turn on to Advanced View, when you go in to the Properties of your Organizational Unit or anything else, you're going to have a lot more tabs. The tab we're interested in is the Object tab. This is where the "Protect object from accidental deletion" is.

Now I can uncheck it and I will be able to delete my Organizational Unit. I like to go back up under View and turn off Advanced Features, just to remind myself I don't need to be in those extra containers. That's how you create Organizational Units. Again, you're going create them to reflect how you plan to support your users, or how you plan to design your group policy.

4.4.3 Organizational Unit Facts

When Active Directory is installed, the following containers and OU are created by default:

- The *Domain container*, which is the root container to the hierarchy.
- The *Builtin container*, which holds the default service administrator accounts.
- The *Users container*, which contains the domain's predefined users and groups. The Users container is also the default location for new user accounts and groups created in the domain.
- The *Computers container*, which is the default location for new computer accounts created in the domain.
- The *Domain Controllers OU*, which is the default location for domain controllers computer accounts.

The default containers are used by the operating system. They cannot be renamed, deleted, or have Group Policy applied to them.

An Organizational Unit (OU) is similar to a folder that subdivides and organizes network resources within a domain.

- An OU can contain other OUs and any type of object type, such as users, computers, and printers.
- OUs can be nested to logically organize network resources.
  - *Parent* OUs are OUs that contain other OUs.
  - *Child* OUs are OUs within other OUs.
  - The recommended maximum nested level of OU containers is five.
  - Too many levels of nested OUs can slow resource requests and complicate group policy application.
- OUs are typically organized by the following:
  - Physical location, such as a country or city
  - Organizational structure, such as the HR, Sales, and IT departments
  - Object type, such as user accounts or computers
  - Hybrid of location, organizational structure, and object type

Be aware of the following considerations for managing OUs:

Feature	Description
Group Policy	One of the main reasons to use OUs to contain objects instead of containers is the application of Group Policy. Create OUs for each group of objects that need to have different Group Policy settings. Keep in mind: • Group Policy objects (GPOs) can be applied to OUs. • Policy settings apply to all objects within the OU. • Through *inheritance*, settings applied to the domain or parent OUs apply to all child OUs and objects within those OUs. A *generic container*, a container created by default, is not an OU and cannot have group policy objects assigned to it. A good practice is to move objects out of generic containers and into an OU. For example, you can move computers out of the Computers container and into an OU, where group policy can be applied.

Preventing accidental deletion	Objects in Active Directory can be accidentally deleted using Active Directory Users and Computers and other management tools. The following types of deletions are most common: • *Leaf-node deletion* is when a user selects and deletes a leaf object. A *leaf objec*t is an object without a child object, also referred to as a *subordinate object*. • *Organizational Unit (OU) deletion* is when a user selects and deletes an OU. Deleting the OU deletes all objects within the OU (including any child OUs and their objects). When you create an OU using Active Directory Users and Computers, the **Protect container from accidental deletion** option is selected by default. You can turn the option on or off after the OU is created in one of the following ways: • On the **Object** tab of the OU in Active Directory Users and Computers. Select **Advanced Features** from the **View** menu before opening the **Object** tab. • On the **Security tab**, in Computers or Active Directory Sites and Services.
Delegating authority	Delegating authority is the assignment of administrative tasks--such as resetting passwords or creating new users--to appropriate users and groups. You should set up the OU structure in a way that best facilitates your support plan. Be aware of the following facts about delegating control: • Using the Delegation of Control Wizard or the Authorization Manager console, you can delegate control of any part of an OU or object at any level. • An object-based design allows you to delegate control based on the types of objects in each OU. For example, you can delegate control over specific object types (such as user objects). • A task-based design allows you to delegate control based on the types of administrative tasks that need to be done. Some examples of administrative tasks are: User account management, such as creation and deletion Password management, such as resetting and forcing password changes Group membership and permissions management

4.5 User Accounts

As you study this section, answer the following questions:

- How is a domain user account different than a local user account?
- What is the difference between a disabled, locked out, or expired user account?
- What is the best way to handle a user's account when an employee quits the company and will be replaced by a new employee in the near future?
- What are the recommendations for using a template user account?
- What permissions does a user account created from a template have?
- How should you re-create a user account that was accidentally deleted?

After finishing this section, you should be able to complete the following tasks:

- Create a user account.
- Disable a user account.
- Manage user account passwords.
- Manage user accounts.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Create and Manage User and Computer Accounts.
  - Create User Accounts
  - Manage User Accounts
  - Configure User Account Restrictions

This section covers the following 70-410 exam objective:

- 502 Create and manage Active Directory users and computers.
  This objective may include but is not limited to:
  - Automate the creation of Active Directory accounts
  - Create, copy, configure, and delete users and computers
  - Configure templates
  - Manage inactive and disabled accounts

4.5.1 Users and Computers

Users and Computers
0:00-0:13
I'm going to talk a little bit about user and computer accounts. The main goals as a network administrator is to have centralized security and centralized administration. That's really what the active director database gives us.

User Account
0:14-0:30
For every user in the organization, there will be a user account, and that user account needs to have a unique name and should be placed in an organizational unit according to how it's going to be supported and what group policies it needs.
Same thing for computers.

Computer Account
0:31-0:43
Computers need a computer account, and then they need to be joined to the domain and linked up with that computer account, and the computer account should be placed in the appropriate location within the organizational structure.

Summary
0:44-0:57
For every user in the domain, you need to make a user account. Put it in the right spot. For every computer in the domain, same thing, we're going to make a computer account and put it in the right spot, wherever it needs to go, in order to properly manage it.

4.5.2 User Accounts

User Accounts
0:00-0:17
Every user in our organization needs a user account that is going to represent them in Active Directory so that we can manage the user account, assigning the right permissions and things like that, and also so that they can get logged in and get access to resources within our network.

Security Identification (SID)
0:18-0:41
When you create a user account, each user account is assigned something called a SID, S-I-D, security ID. It's a unique number assigned to that particular account that never changes, like a social security number in the US. That number is going to be used on any object to which that user has permissions or for any rights that they're going to have.

Developing Usernames Policy
0:42-1:32
When you're creating users, you want to think up a policy for how you're going to develop usernames. Some of the rules that you should follow are: the usernames themselves should be fairly short. You don't want them typing in a novel in there. They should be based on some type of a formula, maybe it's first name, last initial, first initial, last name, complete name. Whatever the formula is going to be, it should be standard and it should account for how you're going to deal with people that have similar names.
If I decide to go with first initial, last name and Shad Farrell becomes sfarrell, what are you going to do when my cousin Sean comes to work for the company? He can't also be sfarrell, because these usernames need to be unique, not only within the domain but ideally within in the forest, so you want to have something already worked out for how you're going to deal with that.

Managing User Accounts
1:33-1:42
Beyond naming the user account and creating it in the appropriate OU, managing it becomes a matter of going in and working with the properties of that account, which we'll look at in the demo.

4.5.3 Creating Users

Creating Users
0:00-0:37
This demo is going to show you how to create user accounts. The tool we need to create user accounts is Active Directory Users and Computers. I'm going to go up and get it from the Tools menu, or you can get it from the Start menu.
When you create your user accounts, you want to make sure you create them in the appropriate organizational unit, so I'm going to be working with an organizational unit called Sales. You can either right click and make a new user, or if you prefer buttons, there's just a little User button up here. Make sure you've clicked on the correct organizational unit, and then you hit New User.

Standards for Creating Name for User Accounts
0:38-1:42
You should give it a First name and a Last name. I believe it requires a Full name. When it comes to logon name, you should create some sort of a rule or a standard so that it's very clear what each user should be named. You also want to make sure you account for people that have similar names, so let's say your organization decides, you know, first initial, last name, and you get Shad Farrell, sfarrell. Well, what happens when my cousin Sean comes to work for the company? He's also going to be sfarrell, so how are we going to deal with that? Is it sefarrell? Is it sfarrell2? You just have to know what the system is, so that you have a standard for creating user accounts. It doesn't necessarily matter that much what the rules are. What's important with any kind of a database, and Active Directory is a database, is to be consistent. The more consistent you are, the more you can leverage that database.
We'll do TUser. None of this is case sensitive, but sometimes if you don't put uppercases, the users will complain.

User Principle Name (UPN)
1:43-1:52
Now, TUser plus this @northsim.com together is called the UPN, User Principal Name.

UPN Suffix
1:53-2:34
This @northsim.com is called the UPN suffix. It's the back of the UPN. Occasionally you might have more than one choice, and so you would click the down arrow and pick the UPN suffix that applies to this user. In my case, all I've got is the name of the domain, so I'm good to go.
The user can logon like that TUser@northsim.com, or they'll be able to use the user name, pre- Windows 2000, of NORTHSIM\TUser. Either one of those will work. If somebody's trying to logon and one doesn't work --- let's say, the UPN doesn't work --- have them try the pre-Windows 2000 logon name and see if that works.

Creating a Password
2:35-2:39
Now, I'm going to give them a password that matches the password policy for the domain.

User Must Change Password at Next Logon
2:40-3:12
It's a good practice to leave, "User must change password at next logon" checked, because what you don't want to have happen is somebody says, Well, it's not me that went in and wrote "the boss is funny-looking with big ears". Shad picked an initial password for me and I never changed it. It must have been him; now if I leave that changed, I know whoever uses that account first is going to be forced to pick a password that they know, and there won't be any security problem saying I did it because I knew the password.
We very rarely would say, "User cannot change password".

User Cannot Change Password
3:13-3:34
That would be if we're making an account for some type of an application, or maybe a kiosk machine. That would probably be the best example, a kiosk machine -- like a machine in a library or a lobby or a cafeteria, where we're just going to make a Client1 user account. We don't want whoever's logged on with that to be able to change the password.

Password Never Expires
3:35-3:45
"Password never expires" is used for service accounts. It overrides the domain password policy, and it's a security risk. There are better options that you can choose, but if you have to, you can.

Prestaging
3:46-3:59
Whenever you make something in advance, we call it prestaging, so if TestUser wasn't going to start work -- let's say, for two weeks -- and I've got the order in to create the user account but I know that account won't be active yet.

Disable/Enable User Account
4:00-4:43
Then what I should do is leave the account disabled. That way it's not a security risk during the two weeks before this user actually comes to work. You can see, if I leave the account disabled, the account gets a little downward arrow showing that it's disabled. If I want to enable it, I can simply right click, Enable Account.
Conversely, if this person goes on vacation or sick leave, or I know they're going to be out for a while, I could come in again and disable the account so that it's not a security risk while they're out on whatever leave they're out on. You can see from the right click menu, the other thing we often do with this is Reset their Password. I also use Add to Group fairly often.

Template Accounts
4:44-5:46
That's the basic steps that you go through to create a user. I want to talk about one more thing, which is template accounts. I'm always looking to set up a system that will maintain itself, for a couple of reasons. Number one, I myself don't always remember what I had for breakfast, let alone what my decisions were six months ago, so I'd like to set up a system where I don't have to think or research the documentation. Documentation is very important, but you have to look at it for it to be useful. The second reason is, if somebody comes in temporarily to fill in for me, or I win the lottery, I retire, somebody comes in to replace me, I want them to be able to use the same system that I've been using.
For that purpose, we often make template accounts, and a template account is simply an account in an organizational unit that's set up the way all the other users in that organizational unit should be set up. And that way, when I create a new user in that OU, I can just copy the template account, and I'll be very close to what I need for that user. We're going to go ahead and create a template account.

Creating a Template Account
5:47-6:13
A template account is just a regular user account, so we're going to be going through the same steps we just went through. I like to name it with an underscore, because that way it'll pop up to the top of the list of users. It's not necessary, but it's just kind of a slick move that keeps it at the top and makes it easy to copy. I'm going to leave my template account disabled.

Setting Up a Template Account
6:14-7:32
Now, what we want to do with the template account is go in and get it set up with what a normal user from Sales would need. Particularly important would be to add it to whatever groups a regular Sales user would need to be added to, so we'll just pick one or two. Let's say everybody in Sales is an Account Operator and a Backup operator. In real life, we wouldn't be using these groups. We might have Sales Users or Sales Managers. Whatever they are, we're just picking a couple for example.
You can go in and set up a bunch of things inside the properties of the user account. Not everything gets copied, so we'll set up some information in here so we can see exactly what gets copied. We've set up a certain amount of information. On my General tab I've got my First name, there's no Last name; Display Name, I put in an Office, Telephone number, E-mail. On my Address tab, I've put in as much information as I have about the address. On the Account tab, we've set some Logon Hours. On the Telephones tab, I put a Home telephone number, and on Organization, I've given it a Job Title, a Department, and a Company.

What Gets Copied When Template Account is Used
7:33-8:51
Now, let's take a look at what gets copied when we use our template account. Now, somebody comes in and joins the Sales department. I can right click the Template account and copy it; give this user a name. Since this is a real user, I'm going to uncheck Disabled. Leave your template account disabled so that it's not a security problem.
You can see it's got my Display Name. It didn't take the Telephone number and it didn't take the E-mail. That's specific to particular user accounts. It's not really part of the template. On Address, it grabbed City, State, ZIP Code, Country, but it didn't grab the Street. It copied the Logon Hours of the template account, it didn't take any of the telephone numbers because that's particular to a user, and in Organization, it didn't copy the Job Title, because that again is specific to the user, but it did grab the Department and the Company. Most important, it maintained the groups that this user is a member of.

4.5.4 Managing Users

Managing Users
0:00-0:03
In this video, we're going to talk about managing user accounts.
Active Directory Users and Computers
0:04-0:19
Anything that has to do with user accounts is done in Active Directory Users and Computers. We're going to launch it from the Tools menu, but you could also use the Start menu.
We have a few users here, and we'll go ahead and play with the user with my name for an example.
Reset Users Password
0:20-0:33
If I need to reset this user's password, I can use the right-click menu. I can also use the right-click menu to disable the account if I need to, or if it was disabled, I could use it to enable it. You can see the downward pointing arrows here indicate to me that the accounts are disabled.
Properties of User Accounts
0:34-0:37
We're interested in the Properties of the user account and we want to go through some of the tabs.
Documentation and Standarization
0:38-1:13
Some of these tabs are really just for documentation. Whether or not you choose to fill them out is up to your company. What you should do is create a standard. The key to managing databases is to be consistent, and Active Directory is a database. If I put in the Office, I should always put in the Office, that way, if I want to go ahead and run a search and say, look, I need everybody who's at the New York office, I know for a fact that I will back good results. If I sometimes put in the Office, then it's not going to be any good for searching. If I'm only going to do it haphazardly, I might as well not do it at all.
General, Address, Telephones, and Organization Tabs
1:14-1:25
General, Address, Telephones, and Organization are all really tabs that you can choose to use for documentation or not choose to use for documentation. Completely up to your company.
Account Tab
1:26-1:28
The tab we work most with is the Account tab.
UPN
1:29-1:50
SFarrelplus@northsim.com is my UPN. The @northsim.com is the UPN suffix. Sometimes we may add UPN suffix to the forest because we want to match up the user's logon name with their email address, and if the back half of the email address isn't the same as the domain, we have to add that in. If I needed to change the UPN suffix, I could do that right here.
Logon Hours
1:51-2:33
Logon hours govern when this user is allowed to logon to the domain. My account's been limited, but by default, all logon hours are allowed just like this. It's really up to you what you decide. If you choose to limit logon hours, the philosophy behind that is a lot of people -- if they're going to do something bad or hacking --they might try to do it when nobody else is around, so we don't want them logging into the network after business hours. If it's not part of their work schedule, we don't want them in there. The only problem with that is, I'll almost guarantee you sometime, somewhere, you'll get a call at 3:00 in the morning saying, "we desperately have to get Shad into the network because he has some project and it's due in the next 10 minutes and we need you to fix that". Really, it's whatever the security policy of your company sets up.
Log On To
2:34-3:11
"Log On To" would specify particular computers that this account is allowed to use. I have very rarely seen this used. Generally you would use it if you're going to set up a kiosk account. There's some type of an account -- a library guy that's used for a computer that's in the library. That account is only going to be used on that computer so for a little bit of extra security, I specify that particular computer. Otherwise, it's probably fine to leave it to all computers. That way, if the user's individual computer dies, they'll be able to logon to some other computer without your having to adjust the user account.

Unlock Account
3:12-3:18
If the account was locked, I could hit Unlock Account in here, and when I click OK or Apply, it would become unlocked.
Password Options
3:19-4:00
User must change password at next logon -- usually we leave that checked when we've reset the password or just created the account. User cannot change password would be used for one of those kiosk accounts, so library guy, we don't want anybody using that account and changing its password, and then we can't login as that user on the kiosk machine. Password never expires overrides the domain password policy. It's generally used only for service accounts. Store password using reversible encryption, does that even sound good? No, this is an older weaker encryption. It's used for backwards compatibility with applications that require NT 40, so you've got some legacy things going on in your network if you need reversible encryption.
Account Is Disabled
4:01-4:20
"Account is disabled" would be used if the account was being prestaged or made in advance, and it's going to be a little while before the user actually uses it; or, if somebody goes out on leave, we want to disable it so it's not a security risk while they're on leave, whether it's a long vacation or a leave of absence for disability, whatever that is.
Account Expires
4:21-5:05
"Account expires" is usually set to Never. We would set an expiration date if we were creating an account for a temporary employee. The manager says, I need an account, but this person's only going to be with us through January 1, 2014. The only tricky thing about this is, it won't warn you when that expiration date is coming up, so what I used to do when I worked in the field was just go into Outlook and set a reminder two or three days before it expires to call the manager of that department, because half the time they'll say, "you said Shad was leaving on the first of January". He says, "no, no, no, we've extended the contract, and then you want to extend that date out; otherwise, on this expiration date, the user will get a message saying, "your account is expired, you cannot log in", and then you have to go in and fix that.
Profile Tab
5:06-5:10
The Profile tab is used to set up logon scripts.
Roaming Profiles
5:11-6:29
If we're going to use roaming profiles, which is a profile that's kept on the server and will be pulled down on whatever computer this user logs into, we would do it in the profile path. We use roaming profiles so that the user's environment will follow them around. The boss comes in and says, "I have a picture of my kids on my desktop, and it doesn't matter what computer I log into in the company I need to see that picture". I would set up a roaming profile. You might think, that's great, why don't we use these all the time. Those profiles, by default, include the Documents folder, and those folders can get very, very big. I can set up a Group Policy that would exclude Documents, but regardless, if I have roaming profiles, I am increasing traffic on my network. Home folder is if I provide a home folder for the user. Many companies will come in and connect a particular letter to a spot up on the server. Here's a little trick: if you do have a server and a share, and you have the appropriate permissions at the time you modify this -- whatever that server and share is -- you can put a %username%, and when I hit Apply or I hit OK, as long as I have rights to that share, the computer will go out and create a folder in that share with the exact same name as the user name, and grant this user the appropriate rights to that share. That's a little trick that you can use to save yourself some time.
Member Of Tab
6:30-6:42
Member Of is an important tab for us because it's how we can go through and add groups to this user account. Whatever groups the user's a member of, they'll get all the rights that pertain to those groups.
Dial In Tab
6:43-5:50
Home Folder
5:51-7:12
Dial In tab is used only for remote access. In this day and age, it's probably VPN. If I allow or deny access in this tab, it's going to override any policies that are set in the MPS network policy, but by default we would prefer to control this via policy. I'm only going to be in here if I have a user that's a member of a group that has access and I want to exclude them by denying it. That's usually the most common scenario where I would be in here in to do this.

Environment And Session

7:13-7:20

Environment is just for remote desktop -- sets up the remote desktop environment. Same thing with Sessions -- controls what's going to happen with the remote desktop session.

Remote Control and Remote Desktop Services Profile

7:21-7:33

Remote Control also has to do with remote desktop, whether or not their session can be remote controlled. The same thing with Remote Desktop Services Profile. Unless you're really using remote desktop, you're not going to mess with those tabs too much.

Summary

7:34-7:42

The most critical tabs that you'll be in all the time would be the Account tab and the Member Of tab. Those are some of the things that you can do to manage the user account.

Summary

8:52-9:21

You can get your template account set up. Anything that's specific to a user account won't get copied. Anything that's generic information will get copied, and most importantly, the group memberships. In that way, once you have your template set up, you can just copy that template, and every new user is very close to where they need to be. You just have to add a little bit of specific information, whatever it is you are adding in your company, and you'll be ready to go. That's how we create a new user, and that's how we use template accounts.

4.5.5 User Account Facts

A user account identifies a single user, such as an employee. Windows has the following types of user accounts:

Type	Description
Local	A *local* user account is created and stored on a local system and is not distributed to any other system. • Local user accounts are created with the Computer Management console. • The local Security Accounts Manager (SAM) manages the user account information. • Only local resources are accessible with local user accounts.
Domain	A *domain* user account is created and centrally managed through Active Directory. • Domain user accounts are created with Active Directory Users and Computers, command-line tools, or PowerShell. • Each domain user account has a unique security identifier (SID) to identify the user. A user can log on to the domain from any computer that is a member of the domain and can access resources on that computer or on other computers for which the domain user account has permissions. • Domain user accounts have a variety of properties, such as user information, group membership, user profiles, and dial-in settings.

Active Directory uses the following name types to recognize each object:

Type	Description
User or Logon Name	The user or logon name is the name of the user account. It is typically a combination of the *given name* (first name) and *surname* (last name) of the user. For example, Andy Waters may have the following logon name, **awaters**. It is best practice to set up a naming convention that identifies how duplicate names will be addressed.
User Principal Name (UPN)	The User Principal Name (UPN) combines the user account name with the DNS domain name. For example, account **awaters** in the **westsim.com** domain would have the UPN **awaters@westsim.com**. • The UPN format is also known as the *SMTP address* format.

	• The DNS domain name in the UPN is known as the *UPN suffix*. • By default, the domain that holds the user account is selected for the UPN suffix. However, you can configure UPN suffixes other than the domain name.
Distinguished names	Distinguished names are the way the Active Directory refers to objects. The distinguished name identifies the full path to an object, including the object name and all parent objects to the root of the domain. The following identifiers are used in a distinguished name: • CN = common name • OU = organization unit • DC = domain controller Each component of the path is separated by a comma. Following is the distinguished name of a computer named **Client1**, in the OU named **Desktops**, in the OU named **Sales** in **Northsim.com**. **CN=Client1,OU=Desktops,OU=sales,DC=northsim,DC=com** If the object is in a container, the name contains the container name, identified as CN=, in addition to the common name. For example, a user named **awaters**, in the **Users** container, in the **northsim.com** domain, has the distinguished name: **CN=awaters,CN=Users,DC=northsim,DC=com** Keep in mind the following about distinguished names: • Each object in Active Directory has a unique distinguished name. • If the name of any component contains a space, put a double quotation mark (") at the beginning and end of the name.
Relative Distinguished Name (RDN)	The Relative Distinguished Name (RDN) is used to identify the object within its container. The RDN needs to be unique only within the object's container. In the example above, the RDN is **CN=awaters**.

4.5.6 User Account Management Facts

Keep in mind the following recommendations when working with user accounts:

Action	Description
Create/manage user accounts	Use Active Directory Users and Computers from a domain controller or workstation with Administrative Tools installed to configure domain accounts: • When creating a new user account: Configure an expiration date for temporary user accounts. Once the account is expired, it cannot be used for logon. Disable an account if the user will be gone for an extended period of time. Disabling prevents the account from being used during the user's absence. Enable the account when the user returns. Configure the logon hours for a user account to allow the account to be used only between specific hours. ▪ Logon attempts outside of the specified hours will not be allowed. ▪ By default, users who are currently logged on when the logon hours expire are allowed to continue working. ▪ To log a user off when the permitted logon time expires, you can configure Group Policy settings to log the user off automatically. Configure a list of workstations that a user is allowed to log on to. This restricts the user to only those workstations specified. • Copy an existing user account to create a similar user account. When you copy an account: You will be prompted for a new name and password. Existing account settings and group memberships will be copied to the new account. Permissions will *not* be copied to the new account. • Add a User Principal Name (UPN) suffix to a forest so that the users who join the forest can use a friendly user-logon name that does not match the domain name. • Authenticate a user who logs on with a certificate by mapping the certificate to the user account. • Restore an accidentally deleted user account from backup rather than creating a new one with the same name. Creating a new account with the same name results in a user account with a different SID that will not automatically assume the permissions and memberships of the previously-deleted account. • Use the **Shift** or **Ctrl** key to select multiple users when modifying properties on multiple user accounts at once. Properties such as the logon name or password cannot be modified in this way.

	• Move user accounts to add them to the appropriate OUs. Grouping users within OUs allows you to apply Group Policy settings to multiple users.
Use templates	If you regularly create user accounts with the same settings, you can create a template account. The template account is a normal user account with the settings you need for subsequent accounts. • Copy the template account whenever you need to create a new one. • Disable this account to prevent it from being used for logon. New accounts retain group memberships but not direct permission assignments.
Manage passwords	Keep in mind the following about user passwords: • When creating a new user account or resetting a forgotten password, reset the user account password, and then select **User must change password at next logon**. This forces the user to reset the password immediately following logon, ensuring that the user is the only person who knows the password. • The **User cannot change password** option allows you to maintain control over a Guest, service, or temporary account. For example, many applications use service accounts for performing system tasks. The application must be configured with the user account name and password. If you allow changing the user account password for the service account, you would need to change the password within every application that uses that account. • To reset the user account password, right-click the user object and select **Reset Password**. • An account that has been locked out due to too many incorrect passwords attempts must be *unlocked*. To unlock an account: Go to the **Account** tab in the account object's Properties dialog box, and select the **Unlock Account** box. Use the Reset Password dialog to unlock a user account.
Create a user profile	The user profile tracks user environment settings, such as program-specific settings, user security settings, and desktop settings (including the files, folders, and shortcuts on the desktop). • By default, the profile is stored on the local computer. A profile is created on each computer when a user logs on. • A roaming user profile makes profile settings consistent across computers by saving the profile to a network share. To use a roaming profile, edit the user account properties and specify the profile path.

	▪ To simplify administration, use the **%username%** variable in the **Profile Path**. ▪ Active Directory replaces **%username%** with the user logon name. When the user logs on, profile settings are copied from the network to the local computer. Changes made on the local computer are saved back to the network share.
Deprovision a user	*Deprovisioning* is the process of removing access rights from a user account when the user leaves the organization. • If the user will be replaced by another user, disable the existing account. When the new user starts, rename the account, reset the password, and enable the account. This process preserves all of the permissions and other settings associated with the user. • If the user will not be replaced, you can delete the account. Be sure to reassign any permissions to other users, reassign ownership over files, or delete unnecessary files such as the user profile. After a user account has been deleted, all permissions and memberships that are associated with that user account are permanently deleted. All permissions and memberships must be recreated manually if you want to duplicate a deleted user account. • Many third-party tools exist that can simplify the deprovisioning process. For example, you can delete the user account and automatically reassign permissions or file ownership with a single step. You can also create your own deprovisioning solution through a programming language to synchronize accounts between databases or applications.

4.6 Bulk User Actions

As you study this section, answer the following questions:

- When would you choose **Csvde** over **Ldifde** when managing objects?
- Which tools add the user password to the user account?
- Which tools can you use to create objects in Active Directory?
- Which cmdlets can be used to manage Active Directory objects?
- What is the benefit of *piping* multiple commands?
- What utilities would you use to view the properties of multiple Active Directory objects?
- What is the default action for the **Csvde** command?

Bulk Operations
0:00-0:09
We are going to talk about bulk operations.
What bulk operations really means is either creating or modifying a lot of objects all at the same time.

Reasons for Bulk Operations
0:10-0:54
There are different reasons why we might want to do bulk operations. One of the reasons might be that we're concerned about making sure that changes are made in such a way that it won't impact the production of the users. When you're planning changes and you know about a change in advance, you should always try to plan so that that change can occur in such a way that it doesn't disrupt the productivity of the users. That's important for the company. If the users aren't being productive, then the company is not making money, and then everybody is going to be sad. Certainly you want to go through and make sure and it might be that you have to change a hundred user accounts, that might be a reason why you do bulk operations. Maybe you're just trying to do it because it's a repetitive action and it would be much more efficient to script it out. There's different reasons why we might want to do this.

Choosing a Utility
0:55-1:02
Depending on what type of a bulk operation you are attempting to do is going to govern which utility you would use.

Active Directory Users and Computers
1:03-1:08
Generally, if we're working with users and computers, that's done in Active Directory Users and Computers.

Bulk Changes
1:09-1:34
If you had bulk changes, you could go ahead and use this to make bulk changes. The key is, the bulk changes would have to all be the same change, because what you would do in here is make a filter -- find all the users, let's say -- highlight them all, and then you can make a change. Let's say, change the company name to NorthSim; everybody is going to get the same change. If it was something where they needed a unique change, this would not be a good utility.

Comma Delimited Files (CSVDE)
1:35-2:11
Another bulk utility we are going to look at is CSVDE. CSV or Comma Delimited Files, where it's kind of like, if you think of Excel, maybe I have a computer name in one column and then I have the type of computer in another column.
This would be a text file and everywhere there would normally be a column break there would be a comma. Then each line would be a different computer. Client1, type is a desktop; Client2, type is a laptop, all right, with commas in between each of these. CSVDE is great for two things.

Bulk Creations
2:12-2:17
It can do bulk creation of objects if I use it to import a Comma Delimited File.

Documentation
2:18-2:31
It can also be used for documentation because if I don't import, I can use it to export and create a Comma Delimited File that has records or documentations about everything that's in Active Directory. The key that you would probably be interested in it for is bulk creations.

LDIFDE
2:32-2:36
Another command that we can use is LDIFDE.

Bulk Unique Changes

2:37-3:14

These are command line commands and this would be used for Bulk unique changes. I don't particularly like this utility because I feel like it's almost easier just to go into each object and make the change. Let's say you had to go through and assign a telephone extension to everybody in the company. I could create a file that would do that. The problem is, I'd have to go through and have the name of the object, Shad, the container name, it's in an OU, and it's quite extensive work to set up the file for LDIFDE, but if you get into a situation where that makes sense, LDIFDE is great for bulk unique changes.

Command Line Command (DS Commands)

3:15-3:54

The last two utilities we are going to use, one will be a command line command. So, we will just call them the DS commands and that is because they all start with DS. Anytime you see DS like that in a command, it means Directory Service, and Active Directory is a Directory Service, so it probably means the command has to do with Active Directory. These are all pretty obvious, so there is a DSadd for adding new items. There is a DSmod for modifying existing items. There is a DSRm for removing items. There is a DSmove for moving items, and basically, I can either execute these one at a time from the command line or I can make a script that would run them.

Batch File

3:55-4:16

If I were to make a script to run the DS commands, the easiest way to do it would be to do a batch file. I would create a text document. I would put one DS command on each line in the text document and then I would save that text document with a .bat extension and then I could just double-click that file. All of these commands would run all at once, one after the other and execute what I need to have happen.

Powershell Script

4:17-4:34

The last thing that we can do for bulk commands would be a PowerShell script. Very similar to my batch file, I would just put my PowerShell commands, one on each line. Again, if I need any commands that aren't in generic PowerShell, I'm going to add an import module command to the top of the script to import whatever module I need.

Summary

4:35-5:08

Those are my utilities for going through and making bulk changes -- changes to a lot of objects or creation of a lot of objects all at once. I can use Active Directory Users in Computers if it is a bunch of objects all getting the same change. CSVDE if it is bulk creation. LDIFDE for bulk unique changes. DS Commands if I want to do it from a command line with a batch file or PowerShell commands if I want to do it from a script using PowerShell.
That way, I can get a lot of things done all at once and make sure that I have plenty of time off-hours to do what I need to do and not disrupt my users during production.

Creating Users in Bulk
0:00-0:22
In this video we're going to take a look at bulk management and creation of objects. So first, let's take a look at our Active Directory structure right now, so that we can see the changes as they come up .. and you can see that we have a domain named northsim.com. We have an OU named Production, nothing in it yet, and we do not have any OUs named Sales or HR.

CSVDE
0:23-0:42
The first thing we're going to take a look at is csvde, and csvde allows you to import CSV files for bulk creation of objects. You can also use it to export information about Active Directory, but the key is, it won't do changes. It will just export information or import new objects.

CSV File In Depth
0:43-2:24
I've created a CSV file here that we can take a look at, and it opens up in Notepad. The first line of your CSV file should identify what each spot is. So, our first line says whatever comes first is the objectClass. So you can see in the second line that object is an organizationalUnit. We have a couple groups. We've got four users and then four computers. Whatever's after the first comma is the distinguished name of that object. So our organizational unit is ou=HR. It's in the northsim.com domain. Got to know how to create the distinguished names. That's very important if you're going to do any type of scripting.
There's a comma. The next spot identifies a givenName. So if you look at that first group line, you can see I've got three commas in a row. It's because groups don't have given names in the SN, and next field is surname. They don't have surnames either. So if you have to leave a field blank, you still need the right number of commas, unless those fields are at the end. Then, you don't have to worry about it. So if you look down at our first user line, I've got James Dow. He's in the HROU that was created earlier. After the distinguished name, givenName is the first name James, surname is the last name Dow, we've got his samAccountName which is going to be his logon name, JDow. UPN JDow@northsim.com. and then displayName is "James Dow", I also threw in some computer accounts down here at the bottom, so we're going to be creating four computer accounts, WS1-4 in the HROU. They don't have first or last names, but they do get a samAccountName, which I just synchronized with the name of the computer.

Import
2:25-3:21
So now we've seen the file; let's take a look at how we do this. This is going to be done in the command prompt. And I want to make sure I open my Command Prompt as administrator, so I'm going to right-click it, Run as administrator. My import-csv file is in a folder called C: Bulk. So we're going to start with the command. You need the -i for import. That's really important, no pun intended. So make sure, if you're importing, it's got a -i for import. If you leave that out, by default, the command is an export. So instead of importing, you'll overwrite your file with an export.
-f specifies the filename and we should be good to go. Alright, you're going to have to refresh. I'm just hitting F5. There's my HROU. You can see it created all my users, created all my computers, and it created my group.

Things Csvde is Unable to Do
3:22-3:25
Here's some things csvde can't do.

Cannot Enable
3:26-3:41
First of all, notice everything comes in disabled. It cannot enable anything. The reason that they really can't be enabled is, they don't have passwords. And I can't set a password with csvde either. It's not going to work.

Cannot Add Members to a Group
3:42-4:00
Finally, I can't add members to a group, so if we look at the Members of our HRUsers group, it's empty. We've got to use different commands if we want to do all those things. But that's how csvde works. Make sure you know the -i for import and it's for bulk creation of objects in Active Directory.

DS Commands
4:01-4:05
Next thing I want to take a look at are the DS commands, and I've created a batch file here.

Batch File
4:06-4:39
These are all text files, they just have different extensions; so import.csv is a text file, it's just got a CSV extension. Batch file is just a text file, but instead of ending in .txt, it ends in .bat. I'm going to right click this and edit it. In a batch file, I don't need a header line like we did with csvde. Each line is just a different command.
So if you take a look, this is going to make a Sales OU. It's going to make a group, SalesUsers, then it's going to add four users, four computers. And then this last command is pretty cool.

dsquery
4:40-4:50
dsquery is used for retrieving lists of things from Active Directory.
So what this says is, get me a list of all the users in this Sales OU, which is going to be those four users up above.

Pipe
4:51-5:35
To join commands together, we use a pipe. And to make this up and down line which we call the pipe, I did a shift backslash, so it looks like two lines above each other on the keyboard, but it comes out as a single line like that. When you have a pipe like that, it says take everything that comes out of the first command, and for each item out of that first command, run it through the second command. So what that's going to do is, it's going to retrieve those four users. For each one of those users, it's going to modify a group, (An existing group called SalesUsers) and add that user as a member of the group.
Make sure you know we cannot create groups with members. All we can do is modify an existing group and add a member. So that last command there is going to find all the users in Sales and add them to the SalesUsers group.

pause
5:36-6:35
I added a pause so that it will stop or pause so that we can read the outcome of the batch file. Otherwise, you just double click it, it opens up, it closes; you can't see if it was successful or if it failed. So let's take a look at our DS commands, okay? And we can see a whole line of these things succeeded. If I hadn't put the pause, the windows would just have closed on its own, so that's fine. I don't see anything where it's complaining, and if I go into Active Directory Users and Computers and Refresh, I now have a Sales OU.
Because these users did have a password, they can be enabled. So I don't have any problems there. And then if I go into my group, I can see that all those users were added as members. So that's how our DS commands work. You can type them out individually at the command prompt. We'll do that in a little bit. But if you have a whole bunch of them at once, it's nice to make a batch file.

PowerShell
6:36-7:31
We're going to take a look at using PowerShell to work with our CSV files. So I have a CSV file here and when you're working with PowerShell the line at the top is a little bit different. So in my CSV file that we used with csvde, it's called distinguished name, and it's got the full distinguished name of the object. Here, I've got to start each line with a field called Name, and then the path tells me what container this object is going to be created in. So it's not the full distinguished name like we saw in the CSV file, it's a little bit different.
And then what I've done at the end is, I've added a password column. csvde cannot create passwords. But we can import the CSV file if we set it up correctly, using PowerShell, which actually can import the passwords. It's a very long command. So what I've done is, I've put the PowerShell command in a text file so I can copy it out. So we're going to open up PowerShell.

import-csv
7:32-8:43
And just to show you what the first command does, we're just going to run it by itself, and what import-csv does is just read through each line of the CSV file and put those entries into variables. So whatever was in the first spot went under the name variable. Second spot was a path variable, given name, et cetera.
So I can take that import-csv file and use the pipe, just like I did with DSad, to put it into a PowerShell command that will create new users. So let's go ahead and open up the file that's got that command ... and this is my command here. And I copy that. And whatever's in the clipboard if you right click a PowerShell prompt -- not only does it put it in there, but it runs the command. I know it went off okay because I don't see anything red.
If we take a look at this command. It's importing that CSV file that we just looked at. For each record in that CSV file, I'm going to make a New-ADUser. For the -name of that user, I'm going to use whatever's in the name variable. It's going to be -Enabled.

AccountPassword
8:44-9:33
Now, -AccountPassword is pretty important. Passwords are not stored in plain text. There's a myth I've heard many times that if you could somehow get to the right spot in the registry you could see everybody's password just sitting there. It's not true. They're kept as a hash. A hash is a unique number where, if you make a hash out of something, there's no way there would be two different strings that would have the same hash. So when you set your password in Active Directory, it converts it to a hash and it's stored in Active Directory as a hash.
When you go to log in, whatever you supply as your password is translated into a hash. That's what actually goes out over the network, and then the hash of what you put in is compared to the hash of what's in Active Directory. If they match, great. If they don't match then you don't get in.

Converting Passwords to Hash
9:34-9:56
In order to take whatever's in the password field and use it as the user's password, I have to convert it to that hash. So I have to add this ConvertTo-SecureString. Whatever's in the password field it requires as plain text, because it's plain text now, and I know you have to add the -force, or it doesn't work. But make sure you know you've got to convert it to a SecureString.

Variable Start With $_
9:57-10:17
And also. you can see each of these variables starts with a dollar sign ($), so I wouldn't put anything in quotes. I want to use whatever's in that variable, and all variables start with a dollar sign ($). So that $_ says a variable in this set period; the name of the variable is whatever's in the password field.

Names of Variables are Case Sensitive
10:18-11:06
The names of the variables will be case sensitive. That's why I ran the import-csv command first, so I can look at the case of each of those variables, and when I type the variable into my command it's got to match whatever's going on behind the scenes. -path is whatever the path is; -givenname, et cetera.
In my PowerShell command, I have to list each of the fields that I'm using. So if I just put in -name is whatever's in the name variable, -password is whatever in the password variable. All those other fields wouldn't be processed. I have to type them out individually.
So let's go ahead and take a look in our Production OU and we can see that those four users have come in. Now, ldifde is used for bulk modifications.

Ldifde
11:07-11:31
So right now, we don't have any Members in the ProdUsers group, and we don't have any Members in the HRUsers group either. ldifde -- I would use that for bulk individual changes, and I'm not going to tell you I'm a fan of this, but we'll take a look at the file. In ldifde, each object that you're working with, you specify the distinguished name.

Distinguished Names
11:32-12:28
In here, the distinguished names are really sensitive, so CN has got to be in capitals. If it's not in capitals, it's not going to work. Normally, if there's a space in the distinguished name, we put the whole thing in quotes. This command does not like quotes. So even though James Dow has a space in his common name, we don't put any quotes. It understands that everything on that line is one unit. So what I'm going to be doing to this group HRUsers, the changeType is Modify.
I'm going to be adding a member, and here's a list of the four members that will get added, and then this bottom part does the same thing to ProdUsers; modifies, adds a member, and adds those four users that are in the production OU.

Command Line
12:29-14:25
So let's take a look. This is also done at the command line. We still need an -i for import, -f for file, so, syntax very similar to csvde. And traditionally, we name the ldifde files .ldf. Now it says 2 entries modified successfully, but there really only are two entries. Modification of this group at the top. Modification of the bottom group.
If you're designing an ldifde file, you have to put a dash at the end of each entry and a space so that it knows where one object ends and the next object begins. So, let's take a look what this did. And it's added all those members to HRUsers and added those members to ProdUsers. So the only thing I've got left really are these disabled accounts and computers inside of the HR OU, and for these, I thought we'd use some DS commands. So our first command here, I'm just going to copy and paste it. It is going to find all the users in the HR OU, and then it's going to modify the users to change their password to that password. They're all getting the same password and -disabled to know

which is going to, in effect, enable the user account. I will copy this and it tells me it's succeeded. There are my user accounts, and then this command does exactly the same thing, but for the computer accounts.
Here I don't have to change any passwords, so it's just -disabled is no, that's succeeded, and you can see that now my computer accounts are enabled.

Summary

14:26-14:48
So csvde for bulk creation, DSadd can be put into batch files for bulk creation, bulk modification, bulk deleting. You can do a lot with the DS commands. PowerShell can be used there, individually or in conjunction with a CSV file. ldifde for bulk modification.

Individual Modifications

14:49-16:37
I usually think of that as individual modifications, and the reason is this: If I needed to find all the users in my domain and make a change, that I could do more easily from Active Directory Users and Computers than I could do from a command prompt, and we can take a quick look at that before we go. I'll go back into Active Directory Users and Computers. If everybody needs the same change, all you really need to do is create a New Query, and I want to find everything where the name has a value. And that's going to bring back all the users in my domain. I hit OK, there's every single user, and then if I want to make a change I just click on the first one, hold down the shift key, click on the last one.
If you want to highlight individually, it's CTRL click; so I can CTRLl click, remove the ones I don't want. If I have a bunch of things highlighted together and I right click one of them and go into Properties, I'm going to see all of the properties that they have in common. Whether it's Account Properties, Address. So I can go in and say alright, you know, everybody here is in the Northsim Office and all of those accounts are going to get that change.
So if everybody's getting the same change, your best bet is Active Directory Users and Computers. I would only use ldifde if everybody's getting a unique change, and then really only if I have a short period of time to make the change and I want to prestage it. If I really did want to edit every single user in my domain and put in, let's say, a unique telephone number, for me it's easier with the mouse. But if my boss says, well, you've got a half an hour after work on Friday, I certainly can make an ldifde file and process it that way.
So that's how we manage Active Directory objects in bulk.

Querying Active Directory with PowerShell
0:00-0:05
In this video, we're going to take a look at querying Active Directory with PowerShell and the DS Commands.

DS Commands
0:06-0:28
The DS Commands first. Those are done on the command prompt. I'm going to make sure that I open my command prompt as Administrator and I've actually typed my commands into notepad to make it a little bit easier.
If you're trying to get information out of Active Directory, your two commands are dsquery and dsget. dsquery finds all the Active Directory objects that match your query.

dsquery
0:29-0:42
We're going to run a query for all the user objects that are in the -HR ou. It comes back and gives me all those four objects.

dsget
0:43-1:07
dsget retrieves the properties of an object. Of course, if you're going to do a dsget, you have to specify which object. You have to do that using their distinguished name, which are these things right here.
I can run in query "find everybody that meets this criteria", then I can do a dsget. In this case I'm saying James Dow. I want to see what his office is. And office is right now set to northsim.
Usually, what we're going to do with this is combine them.

pipe
1:08-1:26
When you combine these things you want to put a pipe in between. This is a shift back slash. Whatever comes out of the first command is going to get piped into the second one. It finds all of those four of those users and shows me the office. That's dsget, dsquery.

PowerShell
1:27-1:48
We're also going to take a look at how to do this from PowerShell.
I'll be showing you get-aduser, get-adcomputer, but there is also a get-adobject. Whenever you're retrieving information, it should start with get. If you're modifying information, it should start with set. Anything new would be either New or Add.

get-aduser
1:49-2:31
We're going to get-aduser. We have to put a filter, but I'm using a star (*) shift A, to say I really want to see everything. These are all the users in my Active Directory environment.
Be careful running a command like that on a production server. If you have a thousand users, that's going to come up with a thousand users in this particular command. If you want to put that into a file in any command, whether it's command prompt or PowerShell, I can add a greater than sign (>) and put the name of the file I want to create. That just created a file in the root of C: called test.txt. If I go in here, it's exactly the same thing that we just saw on the screen, but now I could search for whoever I'm looking for, or maybe I'm just trying to create documentation whatever I need to do.

get-adcomputer
2:32-2:55
Let's take a look at another one. We use the get-adcomputer. These are all the computers in my company, and we'll get a little fancier. This one's going to find all the users. Filter is star (*), but I specify a search base to say Look, I'm really only looking for the users in this particular container. And there are my four users that we saw with the DS Commands.
I can do the same thing with get-computer.

get-computer

2:56-3:57

Here, I'm using a little bit different filter. I've put a single quote around my filter. I'm looking for computers where the name is like -- not exactly, but like -- and then I put my criteria in double quotes. So, I'm looking for everything that starts with "Client". I don't care what comes after client. But it should start with Client, and then the star at the end of Client says doesn't matter what comes after that.

When we did all the computers, we saw a lot of other computers in there. These are just the computer accounts that start with Client and it looks like I've got four of them. One through four.

We'll do the same thing with the user. Here, I'm looking for users whose names start with James. Then, I've added at the end that I want to see all of their properties. But there's actually just one user who starts with James, James Dow, but we've got an awful lot of properties that we can look at for that particular user. Maybe that's not too friendly, or maybe you put that into a text file.

The last one we're going to look at adds this pipe with an FT command.

Format as a Table (FT)

3:58-4:35

FT means Format as a Table. Then, I could literally just list out only the properties that I want. So my table is just going to give me their name and their last logon time. All those are blank, because none of these people have logged on, so nothing to report.

We'll do one more, so you can see something a little bit better. That Format as a Table gives me their name and their SAM account name, which is their logon ID. If I were looking for, let's say, all the users whose last logon is since a particular time, I can go ahead and I can do that as well.

Summary

4:36-4:52

That's how we query Active Directory -- either using the dsget, dsquery commands, or from PowerShell, using a command that starts with get. There's aduser, adcomputer, adobject. You can pretty much query for anything using PowerShell. You just have to get used to how you set up your filters, and what the commands are.

4.6.4 Bulk Operation Facts

Use the following tools if you have a large number of objects to create or modify:

Tool	Description
Csvde	The **Csvde** command imports and exports Active Directory objects using a comma-separated values file. • **Csvde** can read existing information from Active Directory (export) or create new objects in Active Directory (import). • **Csvde** does not modify existing objects in Active Directory. • Use **Csvde** to export objects from one Active Directory system and import them into another Active Directory database. You can modify the file before importing the objects into the second Active directory database. • **Csvde** switches include: **-i** to import objects **-e** to export objects **-f** to identify the filename • When using **Csvde**: Be sure to use the -i switch when importing a .csv file. Export is the default. Passwords are not exported. The added user accounts are disabled. You will need to add passwords and enable them.
Ldifde	The **Ldifde** command imports, exports, modifies, and deletes objects in Active Directory using LDAP Data Interchange Format (LDIF) files. • **Ldifde** files include a **changeType** parameter that identifies the action to take using the data: Add Modify Delete • Common uses for **Ldifde** include: Using **Ldifde** to export a set of Active Directory objects, modify various attributes, and then re-import the file to change the attributes. Exporting or importing data that exists on non-Active Directory LDAP directories. • **Ldifde** switches include: **-i** to import objects **-e** to export objects **-f** to identify the filename

	When you export user accounts with **Ldifde**, passwords are not exported. You can change passwords for existing user accounts using an **.ldif** file, but you cannot add new user accounts with passwords set. To export user accounts and import them with a password, use the following process: 1. Export the user accounts. The **unicodePwd** field will be blank. 2. Import the user accounts to create the accounts. The user accounts will be disabled and the user will be forced to change the password at next logon. 3. Modify the **.ldif** file to change the operation to modify existing objects. Add a password for each user account and add entries to enable the account. 4. Run **Ldifde** using the file with the passwords to modify the existing user accounts. The distinguished name is case sensitive in **Ldifde** commands.
PowerShell	Windows PowerShell is a command-line environment designed for automating administration and maintenance. You can use PowerShell *cmdlets* to create and manage Active Directory objects. Cmdlets can execute single commands or large scripts which can import a CSV file and use the information to create new Active Directory users. For example, the **Import-Csv** cmdlet is used to specify a comma-separated values file containing objects to be imported or exported. The output of this command must be piped to another cmdlet to actually perform the desired operation. For example, to import users, the Import-Csv command would be used first to specify the .csv file containing the users to be added. Then the output would be piped to the **New-ADUser** cmdlet to create new Active Directory users. The Import-Csv cmdlet can also be used in conjunction with other PowerShell cmdlets, such as **New-ADObject**. This cmdlet is used to create many types of Active Directory objects, including users, computers, and groups, as well as sites and subnet objects. Be aware of the following: • When importing a CSV file, use the same variable case in the PowerShell command as the variable name in the CSV file. • Right-clicking at a PowerShell prompt copies the contents of the clipboard to the PowerShell prompt. • Use the | symbol to pipe the output of one command to the input of the next command. • When using a csv file to add user accounts with passwords, include the **(ConvertTo-SecureString $_.Password -AsPlainText -force)** cmdlet in the **import-csv** cmdlet.

Ldp	The **Ldp** utility allows you to search for and view the properties of multiple Active Directory objects. It is a GUI-based, Windows Explorer-like utility with a scope pane on the left that is used for navigating through the Active Directory namespace and a details pane on the right that is used for displaying results.
DS commands	Domain Services (DS) command-line tools are built into Windows Server 2008 and later. The Active Directory Directory Service server role or the Active Directory Lightweight Directory Services (AD LDS) server role installs these tools. Keep in mind: • Run DS commands from an elevated command prompt. • Enter the commands in a batch (.bat) file and execute the .bat file, or enter a DS command at a command prompt. • Use **dsadd /?** to display help for DS commands. Use **dsadd object-type /?** to display help for adding the specified object type (OU, user, group, or computer). • Include a password when adding a user account to enable the account. • When adding an object, the **Dsadd** command allows you to specify if the object is in a security group and the object's scope (global, domain, or local). You can use the following DS commands to perform the indicated action in Active Directory: • **Dsacls** displays the Access Control List (ACL) of objects. • **Dsadd** adds objects. • **Dsget** displays specified properties of an object. • **Dsmod** modifies an object. • **Dsquery** queries Active Directory. • **Dsrm** deletes an object.

4.7 Computer Accounts

As you study this section, answer the following questions:

- When should you pre-stage a computer account in an OU?
- What is the benefit of computer account redirection?
- What is the most likely cause of a computer not being able to join the domain after being turned off for several weeks?
- What must you do after resetting a computer account?
- How can you prevent or limit a domain user from joining computers to the domain?
- How can you join a computer to the domain if it does not have a network connection?

After finishing this section, you should be able to complete the following tasks:

- Create computer accounts and manage computer account properties.
- Redirect the computer container.
- Perform an offline domain join.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Active Directory.
  Create and Manage User and Computer Accounts
  - Create Computer Accounts

This section covers the following 70-410 exam objective:

- 502. Create and manage Active Directory users and computers.
  This objective may include but is not limited to:
  Automate the creation of Active Directory Accounts
  Create, copy, configure, and delete users and computers
  Offline domain join

Computer Accounts
0:00-0:11
We are going to talk about computer accounts.
We are going to talk about creating the computer account, joining the domain, joining the domain when you don't have the domain available to you, and then redirecting the computers container.

Creating the Computer Account
0:12-0:42
Every computer in your organization needs to have a computer account in Active Directory. That computer account can exist before the computer joins the domain, or, if you want to, you can join the domain at the computer and provide administrative credentials and have the account created almost simultaneously with joining the domain. Inside that computer account is a password known only to the domain controllers and the computer. The computer account is always going to exist, even if it's milliseconds, before that computer joins the domain.

Prestaging
0:43-1:19
When we make something in advance, we call that prestaging. If you make the computer account in advance, you are going to prestage the computer account.
Prestaging is a great idea because it lets you create the account in the correct organizational unit. If you decide to simply let the computer account get created as you join the domain, at the computer itself, by default, Active Directory will create that computer account in the computers container. Containers are not great because they can't have group policy applied to them, so we are not going to be able to manage computers in the computers container the way we can manage computers in organizational units.

Redirecting the Computers Container to an Organizational Unit
1:20-1:55
If we want to, we can run a command that will redirect the computers container to an organizational unit so that if anybody does join the domain without prestaging the computer account, instead of that computer account being created in the computers container, it will automatically be created in an organizational unit that you choose. You can only redirect it to one organizational unit. Depending on your Active Directory structure, that may or may not be useful, but at least it would ensure that all computers would have some kind of group policy applied to them even if the administrator didn't take the time to prestage the account.

Command for Redirecting
1:56-2:28
Let me go ahead and give you the command for redirecting the computers container. You would need to be in the C:\windows\system32> directory and the command, I'll write it below just to make sure we can see it, is redircmp. Again, you want your command prompt to be here when you run it, and you would give the distinguished name of wherever you're redirecting it. Maybe I'm going to redirect it to an OU called Desktops in a domain called northsim in a domain called com.

Offline Domain Join
2:29-3:44
Another facet of handling computer accounts is called an Offline Domain Join. In a situation like this, the assumption is the computer that you need to have join the domain is not able to contact the domain controller. Now, you might be thinking, why would I want to join a domain if I can't contact the domain controller?
Maybe you have a situation where you have a vendor that you have image the computers for you that you've contracted with Dell or HP. They put your company image on it and they ship them directly to remote offices, branch offices, and home users, where they're not really able to join the domain. In that case, we can use an offline domain join to have that computer be joined to the domain even though it can't contact the domain. That way, when the user receives the computer, it's set up and it's ready to go. This is our offline domain join.
In order to do this you do need some type of a client that is joined to the domain already. Your first step is on that computer that's joined to the domain, it doesn't have to be a server. It could be Windows 7. It could be Windows 8. You're going to run a command called djoin /provision. As a result of this, it's going to make a text file that Microsoft calls a blob file. You'll run djoin /provision. What that does is, it creates the new computer account up here in Active Directory.

Blob File
3:45-4:33
It files away that secret password in the blob file. The blob file is encrypted, so when you open it, it's all gobbledygook. You need to copy that blob file to the new computer.

I'm just going to put copy blob file, however, you want to get it there, via email, using a flash drive, but the blob file is going to go over to the new computer. Now, at that new computer you're going to run djoin. Again, the switch here is /requestodj for offline domain join and you'll use the blob file with that. What that does is take that password, put it into that new client, and then after you reboot that computer, it will be a member of the domain. Of course, they're not going to be able to logon until it really connects up, but the computer itself will be a member of the domain.

Summary

4:34-4:59

With computer accounts, every computer in your organization needs a computer account. If you just join the domain, it's going to end up in the computers container, but you can redirect that using the redircmp command. You want to make sure everything is at Windows Server 2003 or better. Redirect the computer container if you need to. If you need to join the domain when you're not connected, that's your offline domain join with djoin/provision and request ODJ.

That should be pretty much it for computers.

Dealing With Corrupted Computer Accounts

5:00-5:41

The only last think that I would add is that occasionally these computer accounts become corrupt. If the computer account is corrupt, you will get a message. Something along the lines of the Trust relationship between this computer and the primary domain has failed or the computer account is missing or the password is incorrect. Usually this happens when the computer has been shut down for a good number of days. It changes those passwords every 30 days. So, if it's turned off for more than 30 days, that could create the problem. In that case, what you need to do is reset the computer account and then at the computer itself, join a Workgroup, meaning unjoin the domain, join a Workgroup and then rejoin the domain.

That's everything that you need to know about computer accounts.

4.7.2 Redirecting the Computer Container

Redirecting the Computer Container
0:00-0:04
In this video, we're going to take a look at redirecting the computer's container.
Computer Account
0:05-0:24
Every computer that joins the domain has to have a computer account. By default, if I just go to that computer and I join the domain, the computer account will be created in the computer's container. The problem is exactly that, it's a container object, so I cannot apply group policy to a container.
Prestage
0:25-1:11
What most people do is prestage the computer account, so if I wanted all my computer accounts to, let's say, end up in an OU called Desktops, I would create my OU and then I would prestage the computer account by precreating it. If I do this before I join the computer to the domain, when I actually am at the computer and joining the domain, it will sync up with the computer account.
Notice it says right here, The following user or group can join this computer to the domain. If you do need an end user to join a computer to the domain, you must prestage the account.
In our case, we want to make sure that even if a domain admin joins a computer to the domain, it's going to automatically go into the Desktops OU. It's not going to go into the computer's container.
Redirecting the Computer Container
1:12-1:27
We're going to redirect Computers to the Desktops OU. We're going to need to do this from the Command Prompt. Make sure you right click the Command Prompt and Run as administrator. The command to redirect lives in the c:\windows\system32 directory, so we need to change the prompt over to that directory.
redircmp
1:28-1:45
The command is redircmp, and I need to specify the OU to which I'm redirecting, and now my command is ready to go. Redirection was successful.
Testing
1:46-2:32
Now, just to test it, we're going to go over to a member server, join the domain, and observe that the computer account will be created in the Desktops OU. Our member server is named Member1, and you can see right now it's a member of a WORKGROUP. We're going to join it to the domain and observe where the computer account ends up. I'll hit Change.
Now, we don't have to wait for it to reboot, we can go take a look at our computer account. If I refresh my Desktops OU, you can see the Member1 account ended up inside the OU.
That's how we redirect the computer's container, just to ensure that even accounts that are not prestaged end up in an OU, so that they'll have at least some sort of Group Policy applied to them.
Summary
2:33-2:50
If I need to move them after the fact, I can do that, but that way I'm ensured that there won't be any unmanaged computers in my domain.

4.7.3 Performing an Offline Domain Join

Performing an Offline Domain Join
0:00-0:13
In this video, we are going to see how to perform an offline domain join. The first step is to go to a computer that is a member of the domain. Our local server here is DC1 in northsim.com.

djoin /provision Command
0:14-0:43
The next step is to run the djoin command from the command line. Now, I'm going to right-click it, and then choose Run as administrator. What I need to do is create the computer account and save that information to a file that Microsoft calls the blob file. The command to create that computer account is djoin /provision. We will use djoin /provision. The domain is northsim.com. The computer's name will be Member2.

Send blob.txt File
0:44-2:25
I'm going to save the blob file in a folder on my local C: drive that I premade called djoin ... and the name of the file -- I will call it blob.txt. It could be called anything. It doesn't matter what you name the file or where you save it. The important thing is getting it over to the other computer. It has now successfully created my computer account and it's saved it in the blob file blob.txt.
My next step is to get that file over to the computer that is not connected to the domain. In real life, you might email it; you might send it on disk. Something like that. We are just going to go over to that computer, pull the file over, and then disconnect from the network so that you can see the offline domain join. Here I'm at Member2, which is a member of the WORKGROUP. I actually have mapped a drive over to DC1 so that I can copy the blob file to this computer before I disconnect from the network. Now I've copied my blob file into a local folder c:\blob. The blob file is at the computer. What I'm going to do is go ahead and disconnect from the network so that we can see that the offline domain join really is successful.
You can see I have only one network adapter and I've turned it off, so I don't have any connectivity to any type of a network -- certainly not the domain controller. I'm going to open my command prompt as administrator and to join the computer to the network is also djoin, but this time it's requestodj.

Run djoin /requestodj
2:26-3:10
requestodj, request offline domain join, I'm going to load up blob file which is in c:\blob\blob.text file. The path to \windows is the %system%, which just means use whatever folder windows has been installed into, in this case the c:\windows folder. I'm targeting the local operating system. With Server 2012 you could actually do this to a VHD or an offline image. I'm just doing it to whatever version of Windows is running. You can see that it tells me clearly a reboot is required for the changes to be applied.

Reboot
3:11-3:25
Now that our server has rebooted, you can see that it is a member of the domain even though it is not connected to any type of a network.

Summary
3:26-3:58
Make sure with djoin you know the four steps. You run djoin provision on a computer that belongs to the network. This computer could be a server or a workstation, as long as it is Windows Server 2008 R2 or Windows 7, or better. Then we copy the blob text file to the new computer that is going to be joining the domain. We run djoin requestodj and then we reboot it, so four steps. That is how you perform an offline domain join.

4.7.4 Computer Account Facts

A *computer account* is an Active Directory object that identifies a network computer. The account in Active Directory is associated with a specific hardware device. To identify a specific computer, two processes are required:

- Create a computer account in Active Directory.
- Join the computer to the domain.

You can perform these processes in the following ways:

Method	Description
Pre-stage accounts	Pre-stage a computer account to create the computer account in an OU. • When the computer joins the domain, the computer is matched to the pre-staged computer account. • Use this method to control the location of the computer account in Active Directory.
Manual join	From the computer you are adding to the domain, edit the System properties to join the domain. The computer contacts the domain controller and a computer account is created in Active Directory. When you join a domain and create a new computer account using this method, the computer account is added to the Computers built-in folder in Active Directory.
Redirection	Redirection puts computer accounts normally created in the Computers container into a specified OU. To redirect, enter the **redircmp** command and OU name at the command prompt. Make sure you are in the C:\Windows\System32 directory. For example, to redirect a computer to an OU named **Desktops** in **Northsim.com**, enter the following at the command prompt: **redircmp OU=Desktops,DC=northsim,DC=com**

Offline domain join	During the domain join process, the workstation must communicate with a domain controller. In situations where a network connection does not exist, you can use the offline domain join feature to join the computer to the domain. To perform an offline join, use **Djoin** as follows:
	1. Enter **Djoin /provision** on a computer that can communicate with a domain controller (this computer is called the provisioning computer). This process:
	  ○ Creates the computer account
	  ○ Generates a text file, referred to as the *blob* file.
	2. Copy the blob file to the computer that you want to join to the domain.
	  ○ Run **Djoin /requestODJ** to insert the file into the Windows directory.

	○ Reboot the computer to join it to the domain. You can also use an Unattend.xml file and the blob file during installation to join the computer to the domain during the install process. You can run **Djoin** only on a computer running Windows Server 2008 R2 and later or Windows 7 and later. By default, **Djoin** contacts a domain controller running Windows Server 2008 R2 or later, but you can run **Djoin** with the **/downlevel** parameter to communicate with a pre-Windows 2008 R2 domain controller.

Be aware of the following facts about computer accounts and joining a domain:

- The members of the following groups can create a computer account:
  Account Operators
  Domain Admins
  Enterprise Admins
- After a computer account is created, you must join the computer to the domain before the computer receives Group Policy settings or before Active Directory receives workstation-specific information.
- To join a computer to a domain, you must be a member of the Administrators group on the local computer or be given the necessary rights.
- Use the **dsadd** and **netdom** utilities to join a domain from the command line as follows:
  Use **dsadd** to create a computer account.
  Use **netdom** to rename a computer account.
  Use **netdom join** to join a computer to a domain.

Each computer has a password that is automatically generated when the computer joins the domain.

- When the computer boots, this password is used to authenticate the computer to the domain and establish a secure channel between the computer and the domain controller.
- The password is saved on the local computer and in Active Directory. By default, the password is changed automatically every 30 days.
- If the two passwords become unsynchronized, the computer will not be able to connect to the domain. An error indicating that the computer failed to authenticate is generated. This problem will occur if you have turned off the computer for an extended period, rebuilt the computer, or if you are replacing the computer with another one using the same computer account name.
- When computer logon fails, reset the computer account using one of the following methods:
  Run the **netdom reset** command followed by the computer account name and the domain.
  In Active Directory Users and Computers, right-click the computer account and select **Reset Account**.

After resetting the computer account, you must rejoin the computer to the domain.

4.8 Groups

As you study this section, answer the following questions:

- What are the advantages of using groups when setting permissions?
- What is the difference between a security group and a distribution group?
- What type of objects can be made members of a universal group? A domain local group?
- What happens to user accounts when the group they are in is deleted?
- Which PowerShell commands can be used to manage groups?

After finishing this section, you should be able to complete the following tasks:

- Create groups.
- Create global groups.
- Create a distribution group.
- Change group scope.
- Implement a group strategy.
- Enumerate group membership.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Create and Manage Groups.
  - Create and Manage Global Groups
  - Create and Manage Distribution Groups
  - Change the Group Scope
  - Implement Recommended Group Strategy

This section covers the following 70-410 exam objective:

- 503 Create and manage Active Directory groups and organizational units (OUs).
  This objective may include but is not limited to:
  - Configure group nesting
  - Convert groups including security, distribution, universal, domain local, and domain global
  - Manage group membership using Group Policy
  - Enumerate group membership
  - Manage default Active Directory containers
  - Create, copy, configure, and delete groups and OUs

Groups
0:00-0:52
We're going to talk about groups. Groups are very important. You don't want to address users as individual users. What we want to do is create groups and set up the group structure so that whenever a new user joins the organization, or we have new resources, it's very easy to maintain the system. In order to create a really efficient group system, you need to know the different types of groups and what they're used for. Now, please keep in mind, if you run a single domain, you might not have the need for all these groups. What I'm going to be showing you is the way Microsoft wants you to approach groups.
If you do this, even if you have a single domain, it's going to leave the opportunity open to leverage these different types of groups when your company grows, because, really, you don't see the benefits until you start having multiple domains and then you can see why this is going to save a lot of time and work once you set it up.

It's Better to have More Groups Than Less Groups
0:53-1:23
The other thing I'll say before we talk about them is, I don't know why, but a lot of students are hesitant to create a lot of groups. Nobody is going to charge you per group. It's better to have more groups than less groups, provided you document things well. Good documentation should be a matter of habit with anything that you do in your network.
Let's take a look at the whiteboard and see our different types of groups and how we should use them. First off, we need to start out by deciding whether we need a security group or we need a distribution group. Here's the difference.

Security Group or Distribution Group
1:24-2:18
A security group can do security or distribution.
Microsoft started out saying distribution could be any application, but realistically, distribution means exchange distribution lists. Otherwise, you can choose a distribution group. If it's distribution, it's distribution only. Practically speaking, what does that mean?
It means a security group can be added to the security tab of objects, files, folders, printers. A distribution group cannot. It can only be used by exchange. You might be thinking, why would I ever make distribution group? Well, maybe you need to create a group for the exchange people, but you don't want somebody to use it inadvertently on a resource.
Once you've chosen whether to do a security or a distribution group, then we have to look at the different types of groups available to us. Keep in mind, too, when you're taking a test, make sure you know whether it's security distribution, because that might be the problem with the scenario right there.

Local Groups
2:19-2:25
The first type of groups we're going to talk about are Local groups. Now, you might be familiar with these. These would be in client computers or member servers.

Examples of Local Groups
2:26-3:38
A good example of it would be the Users' group, the Administrators' group, on a client workstation, power users. Local groups can contain Local users.
The only real exception to this is, if they're in a domain, they can contain users from the domain. I try to make this grid as easy to memorize as possible. If you can remember local groups can contain local users, you're in good shape. Where are they used? They're used on the local workstation, so I can use the local administrator's group on that member server or that client only, backup operators are backup operators just for that server.
What they do is they group together rights and permissions. If you think about local groups that you might be familiar with (like users, administrators, backup operators) they group together rights and permissions on that particular client or member server. For example, users groups together all the rights and permissions that I need to use that client or member server.
Backup operators groups together all the rights and permissions to backup anything on that client or that member server, and so on and so forth. They're just used at a local scope. The ones we work with in Active Directory would be the other three.

Domain Local Group
3:39-3:57
Now, Domain Local groups can contain users or groups from anywhere in the Forest, so my Domain or any other domain that's in my forest.
They can only be used within that local domain where they exist. Just like local groups on the computer, they are used to group together rights and permissions.
Examples of Domain Local Groups
3:58-4:08
For example, there's a backup operators group on the domain controllers that groups together all the rights and permissions to back up the domain controllers, and that can contain groups from anywhere in the forest.
Global Groups
4:09-4:21
Global groups can only contain accounts from their own domain, but they can be used anywhere in the forest. This is the one that's kind of different because these are used to group together users and computers.
Examples of Global Groups
4:22-4:47
Some examples of global groups that you might be aware of would be domain users.
Automatically, if you create a user in a domain, it becomes a member of domain users. Domain Admins is used to group together all the user accounts that should be administrators. Domain Computers (every computer that joins the domain automatically) becomes a member of Domain Computers.
The global groups are the ones that are a little bit different because they group together users and computers, not necessarily rights and permissions.
Universal Groups
4:48-5:05
Universal groups, if they're in use, can contain accounts from anywhere in the forest. They can be used anywhere in the forest.
I think of them as grouping together rights and permissions, but really, this column is not applicable to them, and we'll see why in a minute. How do we actually use this information?
Example Domain
5:06-5:10
Let's take a look at a domain. I've created a domain. I've named it northsim.
Active Directory Structure
5:11-5:31
Here's my Active Directory structure. I've got an OU called Sales.
In there, I've got another couple of OUs -- one for Sales Reps, one for Sales Managers. Here, I've got a File Server with a couple of folders on it, one for Data, one for Reports. This is what my Active Directory looks like. Using the different groups allows us to split up the administration.
Creating Global Groups
5:32-5:57
The first thing I'm going to create are some global groups. Remember, global groups group together users and computers. I'm going to do my global groups in red, so I might make a global group called "salesreps_g", so we know that's a global group. Then I'll create another group called "salesmgr_g", which is also a global group, and I will add everybody from this OU into the salesreps_g and everybody from the manager OU into the salesmgr_g.
Creating Global Groups for Security Purposes
5:58-6:51
You might be thinking, why would I create a global group if I already have my users organized into these OUs? Very important to understand that an organizational unit is not a security principle. A security principle is anything that can be added to the security tab of an object, like a file, folder, or printer.
I can't add an OU to the security tab. I need to create groups in order to work with security and permissions. You can also think of it this way. The user account is only going to be in one of my OUs. It can belong to as many groups as I need it to belong to. It's like in real life. I live in one house, but I can be a member of AAA or the YMCA or any type of a gym. I can have lots and lots of memberships even though I only live in one place.
Even if you have your users organized by OU, you still are going to need to create global groups for security purposes.

Network Administrators

6:52-9:17

The network administrators that are handling the users can make sure that these global groups have the correct people in them. There may actually be, in a big company, different administrators handling the servers.

The people that are running the servers will go in and create domain local groups, and I'll put those in blue. I might go in and -- let's just say, for the sake of example, that the managers are going to need modifier rights to any of the folders, and the reps are going to need to Read.

I'll make a domain local group, "sales_m" for modify, and then I'll make another one, "sales_r", for read. I'll give this the modify right to both of these shares, and I'll give this one the read to both of these shares. Now, I have global groups that organize my users, and domain local groups that are organizing my rights and permissions. To give those rights and permissions to the users, I would join them together.

I would take my salesmgr_g and put it into the modify and take my salesreps_g and put it into the read. That way, anybody who's a member of salesreps_g is going to get anything that has been set up for sales_r. You might be thinking, this sounds like an awful lot of groups.

Take a look at why we might want to do this. Suppose my company adds a domain and they made a new tree. Their new tree is westsim.com. It's very similar to the other tree. I have an OU called Sales and then, in there, I have a sub-OU for employees and supervisors.

I'm going to make some global groups, "salesemp_g", "salessupr_g", and then, over here, I have some shares, contracts, and stats, and I'm going to do my domain local groups. I'll make a sales modify group (sales_m), and a sales read group (sales_r). I'll give them the appropriate rights to their shares. Finally, I will add my global groups into the appropriate domain local groups.

Where this becomes advantageous is when you have to share things between domains. This is what we have now in our forest for our global groups and our domain local groups. Management comes to us and says that the managers in each domain should have rights in both of them.

All I have to do is take my sales managers over here, add them into the sales modify, and do the same thing here. I'll take my sales supervisors and add them into the sales modify. Now, suddenly, anybody who's a manager in either of those domains is going to have rights to all of the individual shares in those domains.

Reasons for Choosing Different Group Types

9:18-10:25

Now, this is where it really gets fun, because some people would say, well, Shad, I only use global groups in my domain. Well, if we had only used global groups, when management comes in and says, these sales managers need to be able to use resources all over the forest, remember, global groups can only contain users from their own domain.

What I would then have to do is go to each and every one of those shares and add in the manager's group from the other domain. We only had two shares for each domain local group, but that's still twice as much work as we did by nesting the groups. Suppose we did another way, we said, well, we won't use any global groups at all. We'll do all domain local.

Domain local can only be used in their own domain. In that situation, I would have to go through into the local domain group and pull in all of the sales manager users from the other domain. Then potentially let's just say there's ten users in each OU. That's ten steps instead of one. By having two different levels of groups, I can nest them and make it very easy to share permissions all over my forest.

Universal Groups

10:26-11:38

The last layer of groups are Universal groups. They come in when you want to create a system that will perpetuate itself. I have my two domains, northsim, westsim. The company comes to me and says, we're going to implement six more domains in the near future and we want all the managers everywhere to be able to have rights.

In that case, what I will do is make a universal group. Then I will go ahead and I will put my global groups into the universal groups, and my universal group into the domain local group. What that means is, when my new domains come in, let's say my first domain is eastsim.com.

That local admin in the new domain is going to make a sales manager global group. They're going to make a sales modified domain local group. Then all they have to do is put their global group in the universal group and add the universal group into their domain local group.

Every manager from every domain in the forest is going to have modify rights to those shares, and all of their managers will have modify rights to every share in the forest. Hopefully, I've convinced you that the different levels of groups are useful.

Acronyms for Memorization

11:39-12:08

There's a couple of acronyms for memorizing this that Microsoft uses.

They say, put accounts into global groups, put global groups into domain local, and give permissions to domain local. If you're using universal, we just add a U in the middle here, so these are accounts into global groups. My global groups go into my universal. My universal goes into my domain local, and then I give my permissions to the domain local.

Group Conversion

12:09-12:32

The last thing I want to tell you about is group conversion. You can go from security to distribution and back again. No problem. For the other scopes, you have a little bit of a different story. Global can go back and forth to universal, and so can domain local.

Every once in a while, somebody will say, well, you need to make a global group into a domain local group and vice versa. Very easy. You make your global group universal, make your universal group domain local.

Summary

12:33-13:12

The groups are very, very important because it allows us to save time and set up a system that's going to perpetuate itself.

Make sure you know the different types of groups, particularly domain local, global, and universal, what types of accounts they can contain, where they can be used ... actually, not a bad chart to memorize. Locals all local. Universal all forest. If you memorize one, either global or domain local, the other one is the opposite.

We create these different layers of groups so that we can nest them together, that we have groups that group users and computers, groups that group rights and permissions, and by bringing them together, we make a system that's flexible enough to accommodate for expansion in the future.

Creating Groups
0:00-0:12
In this demonstration, we're going to take a look at how to create groups. In order to create groups, we need to be in Active Directory Users and Computers. You can open that from the Tools menu or the Start menu, whatever is most convenient.

Where to Put the Groups
0:13-0:45
When you're making groups, you have to decide where to put them. Some people like to put the group in the organizational unit, where the users are that are grouped together. Other people like to make an OU called Groups and put them in there. I would recommend that you make a decision one way or another and be consistent.
The key to doing well with the database is being consistent. It often doesn't matter exactly what the rules are as long as you consistently follow them. I'm going to go ahead and make my group in the Sales container.

Make a New Group
0:46-0:50
I can right click and make a New Group.

Name, Type, Scope
0:51-1:00
The first thing I've got to do is give the group a name. In addition to giving it a name, I've got to choose a type and a scope.

Type
1:01-1:00

Security or Distribution
1:01-1:45
My type choices are Security or Distribution. Distribution groups can only be used for Distribution lists, generally used by email for Distribution lists.
They cannot be used for security. If I choose a Distribution group, and I go to the Properties of a folder or a file, or a printer, and I hit Add on the Security tab, I'm not going to see any of the Distribution groups in the list of possible groups that can be added.
Security groups can be used both for Security and for Distribution lists. You might think, Why would I do a Distribution group? Maybe you want a group that's going to be handled just by the email people, but you don't want anybody inadvertently using it for security.

Scope
1:46-1:49
Once I've decided whether it's Security or Distribution, I should decide the group's scope.

Global Groups
1:50-1:53
Global groups are to group together users and computers.

Domain Local Groups
1:54-1:58
Domain local groups are to group together rights and permissions, and Universal is used when we're trying to set up a system that's going to be maintained throughout a forest of many domains.

Universal
1:59-2:52
Some people like to go through and add something to the name so that they can see that it's a Global group just by looking at it. If that's the case, you want to do that consistently. If you don't want to do that, it's going to show us in the description exactly what type of group it is anyway.
I'm going to go ahead and create the group SalesReps. If I expand Type, I can see right away it's a Global group.
If you prefer to work with buttons, you can come up here. In this particular button here, we'll make a new group in the current container. I'm going to go ahead and make a Domain local group, SalesData, _R for read. It will be a Security group, and I'll hit OK.

Adding Users
2:53-5:02
Once I have my group, there's a couple of ways that I can add users to it. One way is to double click the group, go into the Members tab, and add the user in. Be very careful with the tab. Members are a list of users or other groups that are a member of this group. Member Of is a list of other groups that this group is a member of. Those are two completely different tabs.
If I don't want to go into the Members tab and pull the user in, we could also go into Users, Add that user to a group, and they'll be added to the group. That will show up in the Properties of the user on the Member Of tab. Now this user is a member of SalesReps.
We have our acronym AGDLP. We know we want to put our users into Global groups, put our Global groups into Domain local groups, and give permissions to the Domain local groups. I would come in to the Global group "SalesRep". I can do it by saying it's now going to be a member of the Domain local or I could go into the Domain local and click on Members either way.
I'm going to add this SalesData. I didn't type the whole group name, but it will work because it's the only one like that. Now, my Global group is a member of my Domain local group, and I would give rights to the Domain local group. You can see in the Properties of my Domain local group, under Members, there's my Global group.
Often, a couple of different ways to do this; whatever way works for you is fine. Just make sure you know whether you're on Members or Member Of tab. A lot of students will come to me and say, Shad, I can't add this person to a group. They're in another domain. I'm trying to add it to a group. It's because you're on the Member Of tab instead of the Members tab. So you need to make sure about that.
That's how we create groups.

4.8.3 Group Facts

A *group* is used to collect user accounts, computer accounts, and other group accounts into manageable units. Working with groups instead of individual user accounts helps simplify network maintenance and administration. For instance, users in a group receive all the user rights assigned to the group and all the permissions assigned to the group on any shared resources.

Like user accounts, there are both local and domain groups.

- Local groups exist only on the local computer and control access to local resources.
- Domain groups exist in Active Directory and can be used to control access to domain and local resources. Enterprise environments primarily use domain groups to implement user management.

In addition to the group scope, there are two types of groups:

Group Type	Description
Security	A *security* group is one that can be used to manage rights and permissions. • Group members receive the permissions that are granted to the group. • A security group represents an object with a security identifier (SID). Through the *member* attribute, other objects such as users, computers, contacts, and other groups become members of the security group. • A security group can be added to the **Security** tab of an object.
Distribution	A *distribution* group is used to maintain a list of users and is typically used for sending e-mails to all group members. Distribution groups cannot be used for assigning permissions.

Active Directory groups have a group *scope*. The scope defines the potential group membership and the resource access that can be controlled through the group. The following table lists the different security group scopes and their membership and use.

Group scope	Membership	Resource Access
Local	Local groups can contain members only from the local computers or member servers. These include: • The local users group. • The local administrators group.	Local groups can be assigned permissions on the local client to group together rights and permissions. Create local groups representative of the resources to which you want to control

	• A local operators group.	access, and then assign permissions on the resource to the group.
Domain Local	Domain local groups can contain members from any domain in the forest. These include: • Domain local groups in the same domain. • Global groups within the forest. • Universal groups within the forest. • Users and computers within the forest.	Domain local groups can be assigned permissions within a domain. Create domain local groups representative of the domain controller resources to which you want to control access, and then assign permissions on the resource to the group.
Global	Global groups can contain members within the same domain. These include: • Global groups in the same domain. • Users and computers within the same domain. Use global groups to group users and computers within the domain who have similar access needs.	Global groups can be assigned permissions to resources anywhere in the forest. Create global groups to organize users (e.g., Sales or Development).
Universal	Universal groups can contain members from any domain in the forest. These include: • Universal groups within the forest. • Global groups within the forest. • Users and computers within the forest.	Universal groups can be assigned permissions to resources anywhere in the forest. Universal group membership should be relatively stable. For this reason, you should add only global or other universal groups to universal groups. Avoid adding user accounts directly to universal groups.

Group membership is an attribute of a group. To query for group members:

- Using a DS command, enter **dsget** at the command line and specify the group name and other parameters. For example, to get the full name of each member of the **salesusers** group in the **sales** OU in **Northsim.com**, enter:

dsget group "cn=salesusers,ou=sales,cd=northsim,dc=com" -members -expand

- Using PowerShell, enter **get-adgroupmember** and the name of the group. For example, to get the full name of each member of the **salesusers** group in the **sales** OU in **Northsim.com**, enter:

  get-adgroupmember salesusers

  To get the users' names and their SAM account names piped into a table format, enter:

  get-adgroupmember salesusers | FT name,samaccountname

Be aware of the following when managing groups:

- The basic best practices for user and group security are:
  - Create groups based on user access needs.
  - Assign user accounts to the appropriate groups.
  - Assign permissions to each group based on the resource needs of the users in the group and the security needs of your network.
- After creating a group, you may need to convert the group's scope and/or type.
  - Converting a security group to a distribution group removes permissions assigned to the group. This could prevent or allow unwanted access.
  - You cannot directly convert a group from global to domain local or domain local to global. Instead, convert the group to a universal group and apply the changes, then convert the group to the desired scope.
  - If a global group is nested in another global group, the nested global group cannot be converted to a universal group because a universal group cannot be a member of a global group.
- To add or remove members of a group, use one of the following methods:
  - On the group object, edit the **Members** tab and add the group members. Use this method to efficiently add multiple members to the same group.
  - On the user account, edit the **Members Of** tab and select the group to which you want to add the user. The **Member Of** tab displays all of groups to which the object is a member. Use this method to efficiently add a single user to multiple groups.

Because a group can be a member of another group, a group object also has a **Member Of** tab. Adding objects to the **Member Of** tab for a group makes the group a member of another group (it does not add members to the group).

- When you delete a group, all information about the group (including any permissions assigned to the group) is deleted. User accounts, however, are not deleted. They are simply no longer associated with the group. If you delete the group, use one of the following strategies to recover it:
  - Recreate the group, add all the original group members, and reassign any permissions granted to the group.
  - Restore the group from a recent backup.
- Microsoft recommends using the **AGDLP** strategy for creating and managing groups. The goal of AGDLP is to create role-based access controls using nested groups. The AGDLP strategy is composed of the following:
  - **A**ccounts (including both users and computers)
  - **G**lobal groups (representing specific job functions or roles in an organization)
  - **D**omain **L**ocal groups (used to define resource permissions or user rights assignments)
  - **P**ermissions

To implement an AGDLP strategy:

5. Create the necessary user and/or computer accounts.
6. Create a global group and add the accounts as members.
7. Create a Domain Local group in the domain that contains the resource you need to grant the accounts access to.
8. Add the global group as a member of this Domain Local group.
9. Assign permissions for the resource to the domain local group.

4.8.8 Enumerating Group Membership

Enumerating Group Membership
0:00-0:05
In this video, we are going to take a look at how we can see what users are members of a group, using the command line.

Using the Command Prompt to Enumerate Group Membership
0:06-0:25
First, we want to open up the command prompt. I'm going to make sure I Run as administrator. I probably don't need to because we are not going to make any changes, but that's okay. The members of a group are really kept in a property of that group.

DS Commands
0:26-1:03
When you are thinking about your DS commands, make sure you understand dsquery retrieves a list back of Active Directory objects. If I'm looking for any of the properties of those objects, I have got to use dsget. Now I put the distinguished name of the group.
We are interested in the members and I'm going to add -expand, so it shows me their full names. We have four members in this particular group, and there they are.

Using PowerShell to Enumerate Group Membership
1:04-1:19
We can do the same thing using PowerShell. If I'm in a command prompt, I can just type PowerShell to go into PowerShell, or I could open up a PowerShell window. whatever is more convenient. I can tell I'm in PowerShell because I've got a little PS here.

PowerShell Commands
1:20-1:38
Anytime I want to retrieve information, it is going to start with a get, so I'm going to get -- I'm going to put the name of the group, and that should be good enough. There we can see it's our same four users. We got a little bit more information about them.

Adding a Pipe
1:39-1:58
We could actually go through and customize what we want to see by adding a pipe, which is a shift/backlash format as a table that we will look for, so if you want to see their name and their SAMAccountName -- and it looks just like that.
That's how we can check the membership of a group using the command line.

4.9 Rights Delegation

As you study this section, answer the following questions:

- What security principle should be applied when you delegate administrative authority?
- What are the processes typically involved in delegating administrative authority?
- What is a limitation of the Delegation of Control Wizard?
- What are the steps in delegating the right to create and link Group Policy Objects?
- When is it necessary to delegate **Manage Group Policy links**?

After finishing this section, you should be able to complete the following tasks:

- Create security groups and delegate authority based on role.
- Use the Delegation of Control wizard to assign permissions.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 3.0 Manage Active Directory.
  Delegate Administrative Control

This section covers the following 70-410 exam objective:

- 503 Create and manage Active Directory groups and organizational units (OUs).
  This objective may include but is not limited to:
  Delegate the creation and management of Active Directory objects

4.9.1 Rights Delegation

Rights Delegation
0:00-0:22
Let's talk about delegation of rights. One of the fundamental principles that we're going to use to design our OU structure is to keep in mind how we plans to support the users. We want to go ahead and make organizational units such that I can go in and create groups for my network administrators and give them rights to that organizational unit. That way, they can perform whatever they need to inside the OU.

Delegation of Rights Wizard
0:23-1:18
Whenever you have a situation where a group or a user needs to be able to perform functions within Active Directory, you use the Delegation of Rights wizard. You simply right-click the object, whether it's the OU or the domain, and you click Delegate. You can delegate all kinds of things. You can say, Hey, this group is just going to have the right to reset passwords and force them to change their password at the next logon, or, this particular group is going to have the rights to create and link group policy.
Wherever you run the delegation control wizard, those rights will flow downwards. If I have an OU that has two OUs inside of it, if I delegate control at the parent OU, those same permissions will be in play at the child OUs. The only exception to this is when I'm going to delegate out the right to create and link Group Policy Object. Every other right in Active Directory is simply a matter of running the wizard.

Group Policy Object
1:19-1:58
If we're talking about Group Policy Objects, it's actually a two-step procedure. The first step is exactly like you would for any other right. You run your Delegation of control wizard. You may run it wherever you want them to be able to link Group Policy Objects. It could be at the domain, it could be at the OU, wherever that is.
With group policies though, they actually live in a Group Policy Objects Container. The second step you have to do is give them rights to that container. There's two ways to do that. You could Grant rights to the GPO container or there is a group that already has rights. I can Add the users to the Group Policy Creator Owner group. Either of those would be fine.

Summary
1:59-2:20
We're creating an efficient organizational unit structure so that we can divide up the rights, for managing our network. We're going to do that by running the Delegation of Control wizard wherever we want to grant rights, and we'll see in the demo the types of rights that we can give out.
The only time it's more than the Delegation of Control wizard is if it's group policy, and then I've got to do something to get them rights to that Group Policy Objects container.

4.9.2 Using the Delegation of Control Wizard

Using the Delegation of Control Wizard
0:00-0:10
In this video, we're going to take a look at the Delegation of Control Wizard. We want to go ahead and design our organizational unit structure to mirror how we plan to support our users.

Design of Active Directory
0:11-0:34
Let's take a look at how Active Directory's designed in this particular server.
We want to go to into Active Directory Users and Computers. You can see that I have an OU named Sales. There are couple of users in here; Admin1, Sales User1, and the SalesAdmins Security Group.
I also have an OU named HR with an HR, User1.

Delegating Control
0:35-1:12
Let's assume that the SalesAdmins need to support the users in the Sales OU. If I simply added Admin1 to the domain admin's group, that user would have rights all throughout the domain. Maybe I don't want that. I really am just going to limit the members of the Sales Admins group to managing objects inside the Sales OU and not anywhere else in the domain.
You choose the container, whether it's the whole domain or it's an OU, where you're going to delegate control. Right click and choose Delegate Control.

Using the Wizard
1:13-1:19
Welcome to the wizard. Next.
Now, we need to add in the user or groups that we're delegating control to.

Delegate Control to a Group
1:20-1:38
It's better to delegate control to a group. That way, if Admin1 leaves the company, he hits the lottery and retires, just pop that user out of the group. The next user that comes into the company can be added into the group and I don't have to rerun the wizard.
We'll give rights to the Sales Admin's group.

Choosing Rights to Delegate
1:39-1:56
Now we need to choose what rights we're going to give them. You can see there are awful lot of rights that I can give out. Create, delete, and manage user accounts, Reset passwords, Read user information, Change memberships of a group; many common tasks.

Group Policy Links
1:57-2:52
The only one that's a little weird is Group Policy links. This simply gives me the right to manage the links themselves. It would not give me the right to create Group Policy, and the reason is this: if we look in Group Policy, we're going to see a Group Policy Objects Container. That's where the policies actually exist.
Running this Delegation of Control wizard gives me rights to the Sales OU. It doesn't give me any rights to the Group Policy Objects Container. If what I'm interested in is delegating out the right to create and manage policy links, that's always a two-step procedure. One step is to run this Delegation of Control wizard. The other step is to give the Group rights to the Group Policy Objects Container. I can do that by going in and directly giving them rights to the container, or there's a group in Active Directory called Group Policy Creator Owners. I can add them to that group.

Granting Full Control
2:53-3:20
If you want to give full control to this group, I would choose custom, and I would say, "This folder, existing objects in this folder, and creation of new objects in this folder". That pretty much covers everything. We're going to grant full control to the SalesAdmins, but just to Sales.
Now, Admin1 is a member of SalesAdmins, and SalesAdmins has full control to Sales.

Test Rights
3:21-4:43
Let's test Admin1's rights. The best way to test Admin1's new rights without having to log out and log back in again is to do a Run As with Active Directory Users and Computers.
I'm going to right click Active Directory Users and Computers and run as a different user. Now, I've opened Active Directory Users and Computers, but it's opened using the Admin1 account.
Let's test Admin1's rights in Sales. You can see that as Admin1 I have been successful in creating a username, "Sales User2" in the Sales OU.
Let's take a look at what happens when I try do something over in HR. You can see right away that my button to create a new user is grayed out. I also have nothing on the new command, because I don't have any rights in HR. Inside of Sales, Admin1 has full rights because of his membership in SalesAdmins. But over in HR, no rights at all.

Summary
4:44-4:58
That's how the Delegation of Control wizard works. We use it so that we can divide up the administrative responsibilities -- making sure that admins have plenty of sufficient rights where they should be performing administrative tasks, and no rights where they should not be performing any.

4.9.3 Rights Delegation Facts

Delegating administrative authority means not only sharing administrative tasks with other users, but also tightly controlling the permissions granted to each administrator. Use the principle of least privilege to assign users, including administrators, the permissions required to do their jobs, but no more.

Delegation of authority can enhance security by:

- Distributing administrative authority to one or more groups with a more narrowly defined set of responsibilities.
- Decentralizing administrative control.
- Distributing control based on security principles.

Delegating administrative authority typically involves the following processes:

1. Identify administrative roles. Each role describes a specific administrative function or job.
2. Identify the users who perform each role. Create groups for each role and add the users as members.
3. Assign permissions to each group to enable group members to perform the tasks defined by the role.

To allow users to manage Active Directory objects and properties, such as creating users or computers, managing group membership, or resetting passwords, use the Delegation of Control wizard. Be aware that:

- The wizard can assign new permissions only. You cannot modify existing permissions using the wizard.
- The rights delegated at an OU will flow to the child OUs.
- Delegating the right to create and link Group Policy Objects (GPOs) is a two-step process.
    1. Run the Delegation of Control wizard at the domain or the OU where the group should be able to link GPOs. Select **Manage Group Policy links** in the tasks to delegate.
    2. Grant the user or group the rights to access the GPO container. You can either:
        - Grant rights to the GPO container.
        - Add users to the Group Policy Creator Owner group.

When assigning permissions to Active Directory objects:

- Assign permissions to the domain or organizational units (OUs) based on the administrative scope. For example, you might create an OU for a department, then assign permissions to that OU to a department manager.
- Assign permissions for specific object types, attributes, and tasks. For example, you might delegate the authority to manage only the password property of user objects in the OU.

5.1 Single-label Names

As you study this section, answer the following questions:

- When would you use the GlobalNames zone?
- Which strategies can you use to provide single-label name resolution for IPv6 hosts?
- Should HOSTS files be used for name resolution in large networks?
- When will a Windows client use LLMNR? What limitations does relying on LLMNR have?
- How does DNS devolution work?
- How can a DNS suffix search list be configured using Group Policy?

After finishing this section, you should be able to complete the following tasks:

- Configure a DNS Suffix Search List.
- Configure DNS Registration.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  - Manage DNS Configuration
    - Configure Search Suffixes

5.1.1 Single-label Names

Single-label Names
0:00-0:12
We're going to talk about handling what Microsoft calls single-label names. We could call them single-label names. We could say NetBIOS. We could talk about WINS, but basically it's all the same thing.

Original Naming
0:13-1:00
Originally, when the very first Microsoft operating systems came out, there wasn't a standard way of doing names. We didn't even have an internet at that point, and so each computer was given a name like computer1, Client1, server1, DC1, and that was its NetBIOS name. What they would do is send out broadcasts to find out all the other names of the computers on the network, and they would appoint one computer as a master browser. It would keep the browse list, and everybody would register with the browse list and get a copy of the browse list. And we'd be able to find all the computers on the network and everything would be great. If you've been around since the dark ages like me, if you remember back to Windows 95, there might be delay in network neighborhood when you start up a computer until you see it show up there, and that had to do with NetBIOS.

Legacy Applications
1:01-1:40
Theoretically, with Windows Server 2000, NetBIOS went out, and we officially brought in DNS as the name resolution system for Active Directory. The problem is, legacy applications and legacy is simply a politically correct term for old junk. There were a lot of applications that people spent big money on in NT40 that they just didn't want to get rid of when Windows 2000 came out. Everything still worked, no reason to upgrade, so these applications were still dependent on the NetBIOS names. They didn't know how to deal with DNS, and there were a bunch of things that we could go through to try to resolve that.

WINS
1:41-2:46
One of the things that they did was create a system called WINS, and basically what WINS is, is a centralized database for NetBIOS names. The only problem with it is, it's a flat database. Not like DNS, which is distributed hierarchical. So there's a practical limit to how may names any particular WINS server could have. And then you start getting into fault tolerance and it becomes difficult to set up. We could also go in and tell DNS that it could talk to WINS. There were a lot of interim measures that they went through to continue to provide support for NetBIOS, but still try to phase it out.
Now we're kind of coming towards the end of that process, and we're going to take a look at the latest technology, or continuing to support the hopefully very few legacy computers that exist in the network while still pushing everything towards DNS.
So let's take a look at some examples of this. All right, so these are going to be NetBIOS names, or actually, NetBIOS equivalent names. You might see NetBIOS. You might see WINS, which was our centralized way of managing NetBIOS before we really tried to get rid of it or Single Label Names. All of those terms would be used interchangeably.

Single Label Name
2:47-3:09
Here's an example of a single label name. So if I just do a ping Client1, Client1 is a single label name. In order to help you understand this, I'm going to put up the other types of names we use, which is a fully qualified domain name. So in a single label name, it's just the name of the computer versus a fully qualified domain name, which is the name of the computer and the domain in which it lives. Now, DNS can only deal with fully qualified domain names like this.

Translate Single Label Name into a Fully Qualified Domain Name
3:10-3:22
So when you ping Client1, what we've actually got to do is translate that single-label name into a fully qualified domain name, so that we can run it through DNS and get the answer that we need, and there's a few ways to do this.

Primary DNS Suffix
3:23-4:02
When a computer joins a domain, it actually sets the name of the domain as its primary DNS suffix. That primary DNS suffix can be used to try and make a fully qualified domain name that could then be handed over to DNS or name resolution. With the primary DNS suffix, assuming, let's say, the name of my computer is Client1, it's joined the Northsim.com domain, so its primary DNS suffix is Northsim.com. Whenever you go to ping a single-label name, the computer will put them together and make one fully qualified domain name, Client1.Northsim.com, which can then be sent to DNS for name resolution.

DNS Devolution
4:03-5:37
What happens if Client1 doesn't actually live in the domain as the computer I'm using to ping? Maybe I'm a member of Northsim.com, but Client1 doesn't actually live in that domain. There's a number of ways for dealing with that. The first way we're going to talk about is DNS devolution. DNS devolution says basically this: If my computer is a member of a domain and I attempt to contact another computer using its single-label name, my computer can try my primary DNS suffix and all the DNS suffixes in the same tree, so to speak, as my domain name. Supposing that my client is a member of west.Northsim.com, I ping Client1, and because west.Northsim.com is my primary DNS suffix, my computer first tries to find Client1.west.Northsim.com. If that computer doesn't exist, it will then try everything that's in the same sort of tree as me. So I'm in west inside of Northsim, maybe Client1 is in Northsim. It can automatically try that. And if I were four domains deep, it would try all four of them. So it will go completely up the chain until it finds out, yep, it's not anywhere in this chain, and then it will try something else.
If you have computers that do not support DNS, so let's say Client1 is actually running NT40 for some application that's really old and it doesn't have a DNS name at all, it uses WINS for resolution. And we desperately want to get rid of WINS. We want to transition everything over into DNS.

GlobalNames Zone
5:38-6:00
Windows server 2008 R2 brought in a new technology called GlobalNames Zones, and they were intended to replace WINS. GlobalNames Zones are to help resolve single-label names in your own domain. So if this is Northsim.com, I'm looking at Client1 in Northsim.com. Everybody in my domain is going to talk to DNS to get that.

Create Zone Named GlobalNames
6:01-6:08
I would go ahead and create a zone named GlobalNames, all one word. That's my first step. So the actual zone is named GlobalNames, but then I need to turn on Global Names support by running a command on the DNS server, and all commands for DNS start with dnscmd /config /enableglobalnamessupport space, and then 1 turns it on.

Enable Global Names
6:09-6:34
I need a GlobalNames Zone because if this is a client that doesn't support DNS, it can't register itself with a DNS server. It only knows how to register with a WINS server, and I want to get rid of those.

Disadvantages of GlobalNames
6:35-7:06
The only disadvantage to GlobalNames over WINS is that the LAN administrator has to go in and manually create CNAME records for each of the clients. So if for some reason you do have an extensive WINS environment, this won't necessarily actually be a good replacement for it because it's static. But if you just have a few boxes that rely on NetBIOS names, Global Names is a great option.
Now we want to talk about how to resolve single-label names in other domains. DNS devolution and Global Names help me with single-label names in my domain.

DNS Suffix Search List

7:07-9:05

What if I need to be able to contact single-label names in other domains? That's done with the DNS suffix search list. We can manually put this in, in the properties of the network adaptor, but that's not optimal. I don't want to visit every machine and do that. So we actually use a Group Policy called DNS Suffix Search List.

In the DNS Suffix Search List, I simply program a list of DNS suffixes that the computer can try for single-label names. This would be single-label names in other domains. The Client1.Northsim, that's my domain. It will automatically try that, but maybe I also want it to try Eastsim.com and Southsim.com. So I would add Eastsim and Southsim to the DNS suffix search list in Group Policy, and then anybody that's subject to that Group Policy will try all three of these domains to turn that single-label name into a good fully qualified domain name.

Let's look at one more scenario where I've got DNS suffixes in play. So here I've drawn two forests; Northsim.com, Eastsim.com. Got a server, server 1, that's actually a member of Northsim.com, so it's registering itself with DNS as S1. Northsim.com, and it's recording that IP address in DNS. For whatever reason, it has a network adaptor that's connected to Eastsim, and what I'd like to have it do is register the IP address of this network adapter in the Eastsim.com domain as S1 with that IP address. By default, it won't do that. It'll only register itself in the domain that it actually belongs to. So what I would do in a scenario like this is, in the Properties of the network card, there is a textbox that says DNS suffix for this connection, and I would go in and put in Eastsim.com. Then, there's another checkbox that says "Use this connection as a DNS suffix in DNS registration". So that would cause it to use Eastsim.com when it's registering for DNS. Of course, I don't want to go to this server's network card here and do that, so I would fill out this textbox using an option in DHCP, and then I would turn on this checkbox using Group Policy.

Summary

9:06-9:46

NetBIOS names, also known as single-label names, are still in play. Many people try a ping server1 as opposed to server1.Northsim.com. By default, the computer will use its primary DNS suffix, which is the domain it belongs to, to try to make a fully qualified domain name and send that out to DNS. By default, it also uses DNS devolution, so it will try any other domains that are parent domains of where this computer lives. If we need to have single-label names resolve for clients in our own domain, we'll make a Global Names Zone and enable Global Names support. If it's single-label names in other domains, then I will go ahead and program the DNS suffix search list with a list of the domains that the computer can try.

5.1.2 Configuring DNS Suffix Search List

Configuring DNS Suffix Search List
0:00-0:16
In this video, we're going to take a look at configuring the DNS suffix search list to provide single-label name support for single-label name hosts in another domain. These would be NT40 clients that only support netBIOS or work station computers.

Name Resolution
0:17-1:18
Let's take a look at name resolution right now. I can go ahead and ping a computer named member3.eastsim.com. If I just ping a single-label name, I can see that it's adding .northsim.com to the single-label name, to come up with an FQDN. This isn't going to work for member3, because it's not in northsim.com, and my computer can only apply its own domain name; the name of the domain to which it belongs, to that single-label name to try to get a good match. By default, DNS does support DNS devolution, so if I was in west.northsim.com, it would try west.northsim.com, and it would try northsim.com, but if that other domain is not in my branch, I'm not going to be able to ping it.
If I simply try ping member3, I'm not going to get a reply.

Search List Example
1:19-1:53
What we're going to do with our search list is, give the DNS Client a list of domains so that it can try to see if it can come up with a good FQDN and get an answer out of DNS. To set this up, we need to go in to Group Policy. Because I'm actually at a domain controller, I'm going to go ahead and edit the default domain controller's policy, but you should pick whichever policy applies to your clients. We need to expand Policies, Administrative Templates, Network, and then go ahead and click on DNS Client.

DNS Devolution
1:54-2:43
Now we've got some settings for DNS devolution. It's on by default. Once I set up my DNS suffix search list, it's going to take devolution off. If you're counting on that to resolve a bunch of names, you need to make sure you add those names that are in your tree to the DNS suffix search list. We'll go ahead and enable this and then we just add as many suffixes as we need. Then, you can see, it tells us down here, if we need to add a bunch of domains, we would string them together using commas. We will put them in quotes.
Here, I'm just going to add eastsim.com. To get this group policy to take effect immediately, I'm just going to do a gpupdate.
Now that I've added eastsim.com to the DNS suffix search list, my ping member3 should work just fine.

Summary
2:44-3:07
It tried northsim.com, nothing happened, so it went on to the search list, tried eastsim, and got a name out of that, and I got a reply.
That's how we use the DNS suffix search list to support single-label names in other domains.

5.1.3 Single-label Names Facts

On legacy Windows systems, NetBIOS was used to identify computers on the network with single-label names. The WINS centralized database system was created to help resolve single-label NetBIOS names.

In Windows Server 2008 R2 and Windows Server 2012 environments, use the following strategies to provide single-label name resolution:

Strategy	Description
HOSTS file	The HOSTS file is a static file on each client computer that is used for DNS name resolution. While using the HOSTS file on a regular basis requires too much work, you can use it for limited name resolution in the following instances: • To provide single-label name resolution. • To map a hostname to an IP address that is different from what is provided by the DNS server, for example, during testing, when you want to set up an alternate server. • To provide name resolution outside of the local subnet when a DNS server is not used or is not available. Be aware of the following when using the HOSTS file: • Because the file must be configured on each host, its use is impractical except in the case of very small implementations or on a temporary basis. • Because the HOSTS file is checked first in the name resolution process, a query will not be sent to the DNS server if the mapping is found in the HOSTS file.
Link-Local Multicast Name Resolution (LLMNR)	LLMNR is a name resolution protocol that allows clients to find hosts on the local subnet without the use of a DNS server or broadcasts. LLMNR: • Enables hostname-to-IP address and IP address-to-hostname resolution. • Uses multicasts instead of broadcasts. • Operates only on the local link (within a single subnet). • Is used when DNS name resolution fails. • Works with both IPv4 and IPv6. • Is supported on Windows Vista and Server 2008 (and later) and is enabled by default. It can be disabled by adding a registry setting to each client. You can use LLMNR to easily create ad hoc networks, or to find hosts on the local subnet without the use of a DNS server. LLMNR replaces the legacy NetBIOS broadcast capabilities, but requires LLMNR-capable hosts.

DNS suffix search list	A DNS suffix search list is used to locate computers with single-label names in a domain other than the one your computer is in. You can manually designate the suffix name in the properties of the network adapter; however, it is much easier to configure a DNS Suffix Search List using Group Policy. To have a network interface register to a different domain than the domain it is in, use the DHCP option to enter the suffix of the new domain in the **DNS suffix for this connection** text box and check the **Use this connection's DNS suffix in DNS registration** using Group Policy.
DNS devolution	*Devolution* is an Active Directory behavior which allows a client computer from a child namespace to access resources in the parent namespace without the need for a fully qualified domain name (FQDN). The DNS resolver tries to append the parent's DNS name as it goes up the tree until it is successfully resolved or until a specified devolution level is reached. The *devolution level* specifies the number of labels or the size of the parent domain name where devolution will stop. For example, **westsim.com** contains two labels and **corp.westim.com** contains three labels. A devolution level of 3 will use resolution attempts at the **corp.westsim.com** level.
GlobalNames zone	The GlobalNames zone is a special zone in the DNS database that is used for single-label name resolution within a domain. Use the GlobalNames zone to: • Allow clients to use simple host names without domain information for name resolution. For example, to contact a server with the FQDN name **web1.corp.us.westsim.private**, users could simply enter the single-label name **web1**. • Allow DNS clients to contact NetBIOS-only hosts without the need for a WINS server. • Allow IPv6-only hosts to contact NetBIOS hosts (IPv6 does not support the use of WINS).

New DNS client functionality in Windows Server 2012 includes:

- Functionality to reduce traffic over metered links:
  LLMNR outbound queries are not sent to mobile broadband and VPN interfaces.
  NETBIOS outbound queries are not sent to mobile broadband interfaces.
- Functionality to reduce response time:
  LLMNR and NETBIOS queries are sent at the same time and are optimized for IPv4 and IPv6 queries.
  LLMNR and NETBIOS queries are sent at the same time as DNS queries when a network interface is hijacking DNS names.

5.2 Name Resolution

As you study this section, answer the following questions:

- What is the name of the root domain?
- What does the term *distributed database* mean concerning the DNS database?
- How does an *FQDN* identify a host?
- What is the difference between a *forward lookup* and a *reverse lookup*?
- A client sends a request for name resolution to its configured DNS server, but it isn't authoritative for the zone where the name resides. What happens next?
- You've modified an A record to point to a new IP address, but your client still resolves the name to the old IP address. What should you do?

After finishing this section, you should be able to complete the following tasks:

- Configure the DNS service to forward name resolution requests.
- Create a root zone on a specified server.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  - Manage DNS Configuration
    - Configure Forwarders
    - Configure Root Hints
  - Create DNS Zones
    - Create a Root Zone
  - Manage DNS Configuration
    - Configure DNS Registration
    - Configure Forwarders
    - Configure Root Hints

This section covers the following 70-410 exam objective:

- 403 Deploy and configure DNS service.
  This objective may include but is not limited to:
  - Configure forwarders
  - Configure Root Hints
  - Manage DNS cache

5.2.1 Fully Qualified Domain Names

Fully Qualified Domain Names
0:00-0:07
We're going to talk about fully qualified domain names. Fully qualified domain names are the types of names that are resolved by DNS.

Host Names and Fully Qualified Domain Names
0:08-0:18
They're also known as host names, but there's a slight difference there. Host is anything with an IP address that's using DNS. Fully qualified domain name is the full name of that host.

Domain
0:19-0:31
Now, we've got to be careful when we talk about the word domain with DNS, because it doesn't mean the same thing that it means with Active Directory. A domain in Active Directory is different than a domain in DNS.

Example of Fully Qualified Domain Name
0:32-0:35
Here's an example of a fully qualified domain name.

Host
0:36-0:45
Whatever's at the far left, in that first position on the left, is the name of the actual computer or the host. On the far left we have the host.

Root of the DNS Tree
0:46-0:56
On the absolute far right is a period that we normally don't type. It represents the root of the DNS tree, and we'll talk more about the root when we get into DNS name resolution.

Domain: Active Directory and DNS
0:57-1:30
Now, in between my root and my host will be every DNS domain that this client is a part of.
This is what I mean about the word domain being different for DNS than it is for Active Directory. From an Active Directory standpoint, I'm probably looking at two domains: one domain named Northsim.com, and a child domain named west. From a DNS perspective, every time I cross a period, it's a different domain. I have three DNS domains here, .com; Northsim, which is inside of .com, and west, which is inside of Northsim.

Top-Level/First-Level Domains
1:31-1:51
The domain over here on the far right, these are called top-level or first-level domains, and probably we call them that because only people that are actually in the DNS care about the root. You don't actually go in and say "www.yahoo.com". we don't actually have to type that at all. That's just used in the background. This is the top level, or the first level.

Second Level
1:52-1:54
This, then, is the second level.

Subdomains
1:55-2:39
Anything between the second level and the host, even if I had four other words here, they're all called subdomains. I could say third level, but I wouldn't bother. I would call it a subdomain. Here's the reason: If I go out and I register a domain name, the only thing I can register is a second level. I can buy shad.com. I can't register anything below that, so I can't go out and say, hey, I'd like to register shad.microsoft.com. I think that's a great name. microsoft.com owns the second level name, and they can do whatever they want with the subdomains. Beyond second level, we don't normally worry about third, fourth, fifth. It's just a subdomain, and whatever's at the far left is the host.

Summary
2:40-2:56
Fully qualified domain names have the name of the host on the left, and then to the right of that are all of the domains in which that host resides, completely all the way out to the root of the internet. Whenever we write that entire path, we call that the fully qualified domain name.

5.2.2 Internet Name Resolution

Internet Name Resolution
0:00-0:22
We're going to talk about DNS internet name resolution. DNS is the domain name service, but its job is to take a fully qualified domain name and match it up to an IP address. Computers want to use an IP addresses to communicate, but they're not terribly friendly for human beings. I type in a fully qualified domain name, and what I want to get back from DNS is an IP address.
I have a client that needs to find the IP address for www.northsim.com, and we're going to go through the steps that it will take in order to do that.

How a Client Finds an IP Address
0:23-0:32

Check the Cache
0:33-1:03
The first thing my client is going to do is check his cache. A cache is just a little spot in memory where the client stores names and IP addresses that have been resolved in the past. The idea being that, if I'm constantly using the same name, I shouldn't have to go through DNS every time. The entries only live in the DNS cache for a short period of time so that if there are changes and the client hasn't used it in while, it will re-resolve that name. The first place it'll look is the cache, because if it's already done the work in the past, there's no sense in doing it again.

Host File
1:04-2:03
Originally, when the internet was invented, we didn't have DNS. Everything was done with a file called the Host File because there weren't that many computers on the internet. Up to about 1,000 computers, they didn't need anything like DNS. The Host File still exists to this day, and you may sometimes see entries in that cache that come from the Host File. Whatever is in the Host File is automatically preloaded into the cache and stays in the cache unless you remove it from the Host File. The Host File actually exists on every computer in the c:\windows\system32\drivers\etc folder, and it's just simply named Hosts. What I do if I want to add an entry to that Host File is, I need to open an elevated copy of Notepad. I need to run Notepad as an administrator, and then I can add an entry into that file, and what it's going to do is preload the DNS cache.
In general, we try not to use the Host File, because we want to have that centralized name resolution of DNS.

When Do You Use the Host File
2:04-3:25
Here are some examples of where it might come into play. Anytime a computer, one computer, two computers, it doesn't matter; that computer needs a different answer than what's up in DNS. You have to use the Host File. Let's say, for example, I've got an intranet website -- intranet.northsim.com -- and all my clients use that name to get to the intranet website. But I have one web developer that needs to go to a development web server, also named intranet.northsim.com, that has a different IP address and is on a different network. Since that web developer's computer needs a different answer than everybody else in my company, it can't use the answer from DNS, so I would put an entry in the Host File. I'm going to stay away from the Host File. We don't assume that it's in use, and once you put that entry in there, you have to go back and manually change it if the IP address of the server changes.
A good clue that somebody has used the Host File is if you can't remove an entry from the cache. The only reason that entry would stay in the cache and be resistant and not be removed by you is if it's in the Host File. If you'd like to see the DNS client cache, you can just do an ipconfig /displaydns. You can have a space here or no space here. If you need to clear that client cache, you would do an ipconfig /flushdns. If you do that /flushdns and the entry is still there, the Host File is the culprit.

Talk to the DNS Server
3:26-3:54
If the computer doesn't find the entry in the cache, the next thing it's going to do is talk to its DNS server. We'll say that this client has a DNS server named DNS1; that server is also going to go through a set of steps to try to resolve that name, and the first place it's going to look is its cache, because why do the work again if it's already resolved that in the past. If you need to clear the DNS server cache, that's a different story. You would use dnscmd /clearcache.

Authoritative
3:55-4:14
If the entry for www.northsim.com isn't in the cache, the next place the DNS server will look is to see if it's Authoritative for that DNS domain. Am I the DNS server for northsim.com? In our example, DNS1 is not the DNS server, so it's going to keep moving on down the list.

Conditional Forwarding
4:15-5:10
The next thing it will check to see is if conditional forwarding is setup. Conditional forwarding is just what it sounds like. We're going to forward the request or the IP address for this domain name to another DNS server if a condition is met; and the condition would be a particular domain name. Let's say that northsim.com buys all of its computers from a particular vendor; it could be Dell, it could be HP, it doesn't matter. There are a lot of requests for web pages at that vendor's domain. Instead of doing the entire name resolution process through the internet, I get to shoot over to their servers and ask their DNS servers for the answer. We'll go through that more in-depth in another video. Right now we're going to continue past Conditional Forwarding and see what happens next. Conditional forwarding -- forwards if the condition -- a particular domain name is met. In our case, conditional forwarding is not in play, so we're going to keep moving down the list.

Regular Forwarding
5:11-5:50
The fourth thing that my DNS server can try is regular forwarding, and it's just like it sounds like. I'm going to forward the request to resolve this name to another DNS server, but here there's no condition. If I'm not the DNS server, if I'm not authoritative, it's going to somebody else. If I had forwarding setup, it would look something like this: my client will talk to my DNS server, my DNS server checks its cache -- it's not authoritative -- there's no conditional forwarding, not my zone, I pass it off to another DNS server, and then that DNS server is going to start exactly over again. Is it in my cache, am I authoritative, do I have conditional forwarding, do I have forwarding. Eventually, we'll get to someone who gets passed the forwarding step.

Why Would You Set Up Forwarding
5:51-6:42
Why would I want to setup forwarding? Now, a couple of reasons. One reason, it might be that I want to make less work for this DNS server. Maybe I have an agreement with my ISP; they're going to provide DNS and I don't want my server to do a lot of work. Another reason might be that my server is not authoritative for any domain names. A server that's not authoritative for any names is called the Caching Only server, and those servers are just used to speed up DNS access in a spot where there's actually no DNS records being stored. Or I might do it for security; maybe I have internal DNS servers and they're all going to forward to a DNS server in my DMZ, Demilitarized Zone, or Perimeter Network, so that I'm only going to open the external firewall for that DNS server that's in the DMZ. A number of reasons why we might have forwarding setup, but all you really need to remember is with forwarding, I'm completely passing this request on to a different DNS server to do name resolution.

Root Hints
6:43-8:27
If no forwarding is set up, the last thing that my DNS server will use something called Root hints. DNS is a distributed hierarchal database, so there's a hierarchy through the DNS database, and each DNS server is just responsible for a small portion of that database. I'm asking myself, am I authoritative? I'm asking if I have the records for that particular domain.

Once we get into root hints, we're going to go up to the root of the internet, and we're going to begin to follow the trail on that fully qualified domain name from the right to the left until we get all the way to the hosts. You should be familiar with fully qualified domain names; if not, go back and rewatch the video on FQDNs. Let's take a look at what our server is going to do with root hints. Root hints is a list of the root servers on the internet; and a root server is just any DNS server that has a zone literally named ".root", and if you want to, you can even go in and create one on your own server, and you will be a root server. Nobody's going to talk to you, because you won't be listed in root hints, but that's all what root server is.

If I create a .root zone in my server, it will actually go in and gray out Forwarding and Root hints. Because if I'm a root server, I don't need to forward any request to anybody. I'm the top. The head of the internet. I also am not going to need any root hints; I am a root server, I know who I am. If you'd like to see a list and a map of root servers, you can go to www.root-servers.org, and you can see all the root servers out on the internet, and a map of where they're located, who runs them, and their IP addresses. Here's my root server. It's just a .root zone; it's the top of the internet.

Root Hints Scenario

8:28-10:00

Once my DNS server hits root hints, it's going to use root hints to contact the root servers. What it's doing is, it's using this fully qualified domain name as a trail; starting from the far right, at the period that represents the root, and then it's going to find all these different domains. It will use root hints to pick one of the root servers, and it requests information about www.northsim.com.

When clients talk to the DNS server, they call this a Recursive Query. You really don't need that term for your life. It's not going to make or break your DNS Admin experience, but it occasionally shows up in the software and the literature, and what it means is, don't come back until you have the answer. If my client doesn't want any extra work, he just wants DNS1 to give it an answer.

When DNS servers start working using root hints, they use queries that are called "iterative". An Iterative Query means, just tell me what you know. Give me any information that you have, and I'll continue on like a little bloodhound and follow that trail. It's always iterative between DNS servers, unless forwarding is set up where I completely pass the buck to another DNS server. DNS1 contacts the root and asks it what it knows.

The root servers know about the first level servers. The root server knows about .com. When we're going to resolve names on the internet, I always think of that as going up, because eventually somebody's going to go up all the way to the root. When you're working in the field or if you're taking any kind of an exam, a really easy way to remember it is, ask yourself -- you hit a question or a scenario on name resolution, you ask yourself, which way am I going?

Internet Name Resolution (Up)

10:01-10:36

If it's Internet Name Resolution, eventually I'm going to be going up. There's only two ways to go up; Forwarding, where I can completely pass the buck to another DNS server, or Root hints, where I do the work myself. If for some reason internet name resolution is not working, you would need to delete the .root zone on that server, because it thinks it is a root server and it doesn't have forwarding or root hints setup. Once I get to the top, I'm going to start moving down. Anytime we add a name to the left, we are moving down, and the way we move down is through delegation. That's authoritative for the .com domain.

Anything with a sub-domain -- meaning I'm adding a word to the left, maybe I went from northsim.com to west.northsim.com -- I'm going down the DNS tree now, and it's done with a Delegation.

Subdomains (Down)

10:37-12:59

Delegation has two parts; an NS record, which we call the Delegation Record. This identifies the DNS server that's authoritative for that domain. We also have an A Record, which we call the Glue Record, that has the IP address for that server. A delegation is just like what it sounds like in real life. If I delegate something to you, you become responsible for that task. The servers delegate portions of the DNS database to servers below them. Up at the root there's going to be a Delegation. The Delegation Record would say, "if you're looking for .com, there's the NS record, and maybe the .com server's name is dns.com, and that's authoritative for the .com domain".

Now, in order to talk to that server, I need its IP address, so here's where our Glue Record comes in. If you're looking for dns.com, here's its IP address and that information will be sent back to my server. Now my server is kind of like a little bloodhound; it's on the track, so it's going to contact dns.com. Now that it's resolved, it's gotten from the root over here on the fully qualified domain name down to .com. Now it's going to be looking for northsim. Here, because we're moving down, we're still going left, there's going to be another delegation; and I'm just going to abbreviate. Go talk to dns.northsim.com and we'll draw in a dns.northsim.com, whatever the IP address is. That information comes back to my DNS server. My DNS server is getting close, it contacts the dns.northsim.com server. That server has the northsim.com domain, it's authoritative, and it says, hey, www, that's my pal, I absolutely know that thing's IP address. I've got an A Record that says www, it's some particular IP address that goes back to my DNS server, and the first thing my DNS server is going to do is put it in its cache so that we don't have to do this again if somebody asks right away.

Then it sends the answer to my client, client sticks it in the cache, and now that the client's got the IP address, it can contact www.northsim.com using the IP address.

Summary

13:00-14:02

That's how name resolution works on the internet, and even in an intranet, it will do the same thing if you had multiple domains.

Really focus in on the main concepts. First, the clients check their cache, which can be cleared. If any computer needs a different answer than DNS, that goes in the Host File. If I get to the DNS server, it looks in its cache, is it authoritative, is it the DNS server for that domain, conditional forwarding, then regular forwarding, where it just completely passes it off to another server and root hints. Again, if I'm doing internet name resolution, only two ways to do that: forwarding and root hints.

Once I get up to the root, I'm going to start moving down the DNS tree, adding words to the left in the fully qualified domain name. There's only one way to go down, which is a delegation. When you're faced with doing this in real life or any kind of exam, just ask yourself, what direction am I going? Am I doing internet name resolution going up? Then I'll choose forwarding and root hints, or am I or am I going down adding word to the left, and then I'll choose the delegation? That's the main concepts of internet name resolution.

5.2.3 Adding an Entry to the Hosts File

Adding an Entry to the Hosts File
0:00-0:10
In this video, we're going to take a look at adding an entry to the hosts file.
Any time you want one computer to have a different answer than it would get from DNS, you need to use the hosts file.

Opening An Elevated Copy of Notepad
0:11-0:21
The best way to do this is to open an elevated copy of Notepad. I'm going to right-click Notepad and then click Run as Administrator.

Open the Hosts File
0:22-1:08
Now I've got to open the hosts file. The hosts file is kept in the C:\Windows\System32\Driver\etc folder. I don't see anything because right now, Notepad is set to only display test documents and the host file does not have an extension. I need to click the down arrow and say that I would like to see All Files, and then I'm going to open up hosts.
The entries that have the number sign in front of them are just comments. You want to click on the first available line and put in the IP address. You hit Tab and then put in the name of the computer, and then save your hosts file.

DNS Cache
1:09-1:45
In this case, I don't actually have a computer named www.northsim.com with an address of .60, but I can make sure that it's taken effect. Any entry in the hosts file is automatically put into the DNS cache.
I can look at the DNS cache using ipconfig \displaydns. You can see www.northsim.com; there's the IP address. Even if I flush the DNS cache, that entry is not going to go away, and you can see the entry is still there.

Summary
1:46-1:58
Any time you want to have one computer have a different answer than DNS, you use the hosts file. If you can't get rid of an entry from the DNS cache, in that case, I would check the hosts file.
That's how we add an entry to the hosts file.

Configuring Forwarding and Root Hints
0:00-0:12
In this video, we're going to take a look at Configuring, Forwarding, and Root Hints. For this, I want to open up the DNS Console, so I'm going to right-click DC1 and I'm going to hit Properties.

Forwarding
0:13-0:45
Forwarders is the tab that I would use to set up forwarding. When you set up forwarding, you're basically sending all of your requests to another DNS server. Your DNS server is only going to resolve requests for which it is authoritative, meaning it has those zones. Anything that's not a zone living on this server will get kicked over to whoever I'm doing forwarding to. I would just hit Edit and put in the IP address of the computer to which I would be forwarding, wherever that computer is.
Notice, by default, it's going to say "Use root hints" if no forwarders are available.

Use Root Hints If No Forwarders Are Available
0:46-1:04
If there's nothing here for forwarding, then I'm going to hit root hints. If not, it will use the forwarding. Forwarding trumps root hints. Forwarding is set up; that's going to be the rule. We're not going to hit root hints.

Root Hints
1:05-1:30
The only way we're going to hit root hints is if we get all the way down to the bottom of the places that this computer's going to check. It's going to first check its cache and then it's going to check and say, "am I authoritative", meaning, do I host the zone. For example, on this particular computer, it's authoritative for eastsim.com, northsim.com, and this _msdcs, which is really Active Directory in northsim.

Conditional Forwarding
1:31-1:58
If it's not any of those zones, my next step is to go to conditional forwarding. If that's not set up, I'll hit regular forwarding. If this regular forwarders tab is not set up, then it goes to root hints. Root hints are kept in a file on the server and it's pre-populated. You can see that, even though I haven't done anything to my computer, it's got all these different root servers in here. We'll just focus on a 198.41.0.4.

Root Servers: www.root-servers.org
1:59-2:58
If you're interested in taking a look at the root servers, you can go to an address on the internet. I'm going to go ahead and show you that particular address in internet explorer.
The address is www.root-servers.org, and each of these bubbles represents a root server on the internet. If I scroll down, it tells me the addresses of them, and you can see that the root server named A (over here on the right) does in fact have address 198.41.0.4, and you can see who runs the root server -- where they're located.
As soon as I install DNS it's got all these addresses on its Root Hints tab, and it's going to go ahead and try to contact the nearest root server so it can start doing name resolution out on the internet.
One thing to know about forwarding and root hints, if you are a root server yourself. What is a root server?

Root Server and Root Zone
2:59-3:42
Let's take a look at that. It's a server that hosts a copy of the root zone. Well, the root zone is just a zone named with a period. Here, pretty much congratulations, I'm a root server on the internet. If I'm a root server on the internet, there's no need for me to forward to anybody. I'm at the top of the internet. There's also no need for me to have a list of my colleagues. Again, I'm at the top of the internet. I have a copy of this zone. If I go back in to the properties of my server, you see forwarding is grayed out. I can't set up any forwarding. Root Hints has been cleared out, and I can't set that up either.

Deleting a Root Zone
3:43-4:16
If you have a situation where these tabs are grayed out or it just can't be done, maybe you need to come in and delete a root zone.
Here, I've deleted the root zone and it's prompting me to add the root hints back in, and then I can click OK.
There was a period in time where if you just hit Next, Next, Next, Next on the Active Directory install, it would create a root zone, or maybe somebody in your environment has created a root zone because they wanted an internal root. For whatever reason, if those are grayed out, then you've got to delete that zone.

Summary

4:17-4:30

If we're trying to enable computers to have internet name resolution using our DNS server, there's two ways to do it; one, set up forwarding to another DNS server, which I'm doing either because I have a very complex environment or I'm concerned about security, or let it use root hints.

5.2.5 Intranet Name Resolution

Intranet Name Resolution
0:00-0:08
We're going to talk about a little bit different type of name resolution. You could think of it is as an intranet resolution, but really it's a little bit more in-depth than that.

Name Resolution Review
0:09-0:22
Well, let's have a quick review of what we know about name resolution from our internet name resolution video. If we're doing internet name resolution, eventually we're going to have to get up to the top of the DNS tree and talk to the root servers.

Internet (up): Forwarding and Root Hints
0:23-0:36
There's really only two ways to go up and do that internet name resolution, which is forwarding and root hints. If I'm going up to the Internet (up) Forwarding or Root hints. If I don't want any internet name resolution; I'll just turn those off.

Subdomains (down): Delegation
0:37-0:43
Once I get up to the root and I start moving down or adding a domain to the left, the only way that can be done is with a Delegation.

Sideways
0:44-0:54
What we're going to talk about in this video is a completely different direction. I'd like to think of it as Sideways. When I'm going sideways, I've got two domains that need to talk directly to each other.

Scenario One
0:55-2:01
Here, I've developed a company northsim.com and eastsim.com; they both had been excellent in following Microsoft's recommendation, naming their Active Directory domains corp.northsim.com and int.eastsim.com. It's not guaranteed, but you usually know you're going sideways when you see that you're going directly between names. They're often of equal length, although they don't have to be.
The key is this: these domains need to resolve names in each other's domain without going through the internet. One or both of them is not resolvable on the internet. This often happens in the case of a merger. What I want to see happen is, when a client in corp.northsim.com contacts its DNS server, the DNS server cannot go up through the root of the internet. We want it to shoot sideways and talk to the DNS server in eastsim.com that's authoritative for that domain and give the answer directly back. That's one scenario in which we might need to go sideways.
Two domains, separate companies, but one or more of them is not resolvable from the internet.

Scenario Two
2:02-3:05
Another scenario where we might need to go sideways would be if we had a really convoluted Active Directory structure. Let's say my client is down here in A.B.C; and it wants to request information on a resource that's over here in A.D.E. I've just tried to keep the names really simple because we're not worried about what the names actually are. We just want to look at the process. The client will contact its server and say, "Hey, I'm looking for this server over in A.D.E, this DNS server has got to go up to the top of that DNS for the forest, it'll talk to the DNS server in A, the DNS server in A will bounce it down to the DNS server in A.D, and that DNS server will have another delegation that bounces it down to the correct DNS server, and then finally, the answer will come back". That's a pretty convoluted process if we consistently have a lot of activity from here and over there. What I've really like to have happen is for the client to contact its DNS server and for that DNS server to immediately shoot sideways and contact the correct DNS server for that domain, so that the answer can come directly back.

Ways to Move Sideways: Conditional Forwarding and Stub Zone
3:06-3:41
Once you understand why we would want to move sideways, so to speak in the DNS tree, we then need to look at the two ways that we can do that: conditional forwarding, or to use a stub zone. The really great thing about these are that the pros and cons of conditional forwarding are exactly reversed in a stub zone and vice versa. If you're trying to think of a quick way to remember what we're going over, make sure you know what it means to go sideways. When you see a scenario like that, if you know the pros and cons of one of my methods, you should be able to figure out the pros and cons of the other one, and then you'll know which one to choose.

Conditional Forwarding

3:42-5:06

Conditional Forwarding is exactly that. It's going to forward the request to another server based on a condition, the condition being a particular domain name. In the previous example, if that domain name is A.D.E, I'm going to shoot over to the right.

The problem with this, or the disadvantage of this, is that it's static. I'm going to go into DNS and say, here's the IP address to that DNS server. If that should change, I'm going to have to revisit DNS to put in the new IP address. The pros are, first of all, I don't need permission to do conditional forwarding.

In order to have a copy of somebody's zone -- their DNS records -- I need to have permission. With conditional forwarding, I don't need any permission; just go in and say, hey, if it's this particular domain, send it over to this particular IP address. I can do that even with companies out on the web if I find that there's a lot of internet traffic with their DNS servers.

Since I don't have a zone, the other great thing about conditional forwarding is there's no transfer of records. Sometimes we'll see that, where for whatever reason, the requirements of the company say there should be no transfer of records across this link or that type of thing, in that case, we would have to choose conditional forwarding.

My pros and cons for stub zone are exactly the opposite. First, I want to talk about what a stub zone actually is.

Stub Zone

5:07-5:44

A stub zone is a copy of the zone, which is the place where the records for that domain are kept; but it doesn't have all the records for the domain. It only has the records that identify the DNS servers for that particular domain. That means that a stub zone has far less records than a regular copy of the zone, and it's not going to have much replication or zone transfer traffic. The only time records are going to go across the wire to my stub zone are in the event that an IP address of one of the DNS servers changes. Unless a DNS server changes its IP address, there won't be any traffic.

Scenario

5:45-6:56

Let's take a look back at this scenario. Suppose eastsim.com has 5,000 clients and 10 DNS servers. That int.eastsim.com zone is going to have 5,021 records. 5,000 records for the client, each DNS server gets two records, one that identifies has its DNS server, and one that has its IP address, and then there's one record that identifies the zone. The full copy of that zone then would have 5,021 records. A stub zone would just have 21; the 10 records that identified the DNS servers, the 10 records for their IP addresses, and the one record that identifies the zone. The great thing about our stub zone is, its, dynamic. If any of those DNS server IP address change, I will get that information sent over to the stub zone.

Problems with it? I need permission. I'm only going to be able to have a stub zone for DNS domains where I have a relationship with that company. They've got to go in and give me permission. If the company requirements say, "No changes to the zone", then you can't get permission; you've got to go with conditional forwarding. The other negative is, there will be some -- reduced, but still some -- transfer of records. With this type of name resolution, we looked at some reasons why we might need to go sideways.

Summary

6:57-7:37

Sometimes it's in the intranet, a lot of times its between companies. One might be that there's one of the domains, or maybe both of them, that are not resolvable through the internet using root hints, or, it might just be that we have a convoluted Active Directory structure, and we want a shortcut to speed up some of the DNS name resolution queries. Regardless, if I'm going sideways, I only have two ways to do it: conditional forwarding, or make a stub zone. Conditional forwarding is static, but I don't need any permission and there's no transfer of records. Stub zone is awesome because it's dynamic, but unfortunately I'm going to need permission, and there will be some -- though minimal -- transfer of records between the zones.

5.2.6 Name Resolution Facts

The Domain Name System (DNS) is a hierarchical, distributed database that maps logical hostnames to IP addresses. With DNS, users reference computers using logical hostnames, and those hostnames are translated to IP addresses using DNS. A DNS server performs this service on a TCP/IP network. You should know the following facts about DNS:

- A DNS server holds a database of hostnames and their corresponding IP addresses. Clients query the DNS server to get the IP address of a given host.
- The DNS hierarchy is made up of the following components:
  - . (dot) domain (also called the root domain)
  - Top Level Domains (TLDs) (.com, .edu, .gov)
  - Second-level and additional domains
  - Hosts
- A fully qualified domain name (FQDN) includes the hostname and the name of all domains back to root.
- DNS is a distributed database; no one server holds all of the DNS information. Instead, multiple servers hold portions of the data.
  - Each division of the database is held in a zone database file.
  - Zones typically contain one or more domains, although additional servers might hold information for child domains.

Be familiar with the following DNS terms:

Term	Definition
Forward lookup	A *forward lookup* uses the hostname (or the FQDN) to find the IP address.
Reverse lookup	A *reverse lookup* uses the IP address to find the host name (or FQDN).
Authoritative server	An *authoritative* server is a DNS server that has a full, complete copy of all the records for a particular zone.
Referral	*Referral* is the process by which DNS servers use one another to resolve requests from their specific clients. Because each DNS server is responsible for a small piece of the DNS namespace, the servers contact one another when they cannot resolve queries from their own clients. For example, a root DNS server refers DNS servers to .com, .edu, or .gov DNS servers.
Recursion	*Recursion* is the process by which a DNS server or host uses root name servers and subsequent servers to perform name resolution. Most client computers do not perform recursion, rather, they submit a DNS request to the DNS server and wait for a complete response. Many DNS servers will perform recursion.

Iterative	*Iterative* is the process by which a DNS server requests information from other DNS servers and maintains responsibility for resolution.
Delegation	*Delegation* is the process by which a DNS server hands responsibility for the request to another DNS server.

You should be familiar with the DNS name resolution process that occurs on the client and on the server:

Location	Process
Client	The process for DNS name resolution on the client is: 1. Entries in the Hosts file are preloaded into the cache. o The Hosts file is located in the **c:\windows\system 32\drives\etc** folder. o All changes to the Hosts file are made manually. To change the Hosts file, open the file in Notepad with elevated privilege and make the changes. o You should use the Hosts file only when you want to direct the DNS client to a host different than the host listed on the DNS server. o If you cannot remove an entry from DNS cache, check the Hosts file. 2. The client examines its local DNS cache for the IP address. The cache holds in memory hostnames that the client has resolved in the past. Entries stay in memory only a short time so that hostnames in the cache are periodically re-resolved. 3. If the IP address is not in the cache, the client sends the request to the DNS server.
Server	The process for DNS name resolution on the server is: 1. A DNS name resolution request is forwarded to a DNS server. 2. The DNS server examines its local DNS cache for the IP address. To clear DNS server cache, use the **DNScmd /clearcache** command. The DNS server cache is not the same as the client cache. Windows 2008 Server and later has a DNS client cache, but this cache is not used to respond to client requests. 3. If the name is not resolved using the local DNS cache and the DNS server is authoritative, the DNS server responds using information in the zone hosted on the server. If the DNS server is not authoritative, but is configured for forwarding or conditional forwarding, the DNS server forwards the request.

4. If the DNS server cannot forward the request, or if forwarding fails, the DNS server uses its Root Hints file (also known as Cache.dns). The Root Hints file lists the known root DNS servers.

   www.Root-Servers.org lists the root servers and displays a map of where they are located.

5. The root DNS server responds with the address of a com, edu, net, or other DNS server (depending on the request).
6. The DNS server forwards the request to the high-level DNS server, which can respond with a variety of IP addresses.

Keep in mind the following facts regarding root hints and the root zone.

- The root zone is at the top of the DNS hierarchy, and is named **.** (dot).
- The root servers delegate portions of the DNS database to servers below them. Delegation continues downward until the IP address of the host is found.
- If you have a root zone configured on a DNS server, the server will act as a root zone server.
- A DNS server configured as a root zone server will never use the root hints. It considers itself authoritative. Consequently, the server won't access the Internet to forward DNS queries.
- If you want the DNS server to access the Internet, delete the root zone in the DNS console.
- You can configure root hints through the properties of a DNS server.

5.3 Zone Management

As you study this section, answer the following questions:

- Why is a reverse lookup zone for a network written backwards?
- What are the differences between *Standard Primary* zones and *Standard Secondary* zones?
- What is the difference between *Refresh Interval* and *DNS Notify* triggers for zone transfers?
- What information does an SOA record contain?
- How is an Active Directory-integrated zone different from a primary zone?
- What type of zone would you create if you wanted to use secure dynamic updates?
- What type of name resolution is performed by reverse lookup zones?

After finishing this section, you should be able to complete the following tasks:

- Create a new zone and configure it to be stored in Active Directory.
- Manage zone properties.
- Convert an existing zone to an Active Directory-integrated zone.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  Create DNS Zones
  - Create an Active Directory-integrated Zone
  - Convert a Zone to Active Directory-integrated

This section covers the following 70-410 exam objective:

- 403 Deploy and configure DNS service.
  This objective may include but is not limited to:
  Configure Active Directory integration of primary zones

5.3.1 Forward and Reverse Lookup Zones

Forward and Reverse Lookup Zones
0:00-0:04
We're going to talk about forward lookup zones and reverse lookup zones.
Forward Lookup Zones
0:05-0:14
Forward lookup zones resolve names to IP addresses. You have a fully qualified domain name, and you're looking for an IP address. What you want to create is a forward lookup zone.
Reverse Lookup Zones
0:15-0:20
Reverse lookup zones, then, are exactly that -- the reverse. They resolve IP addresses to names.
Fully Qualified Domain Names
0:21-1:21
There's only one other thing that you really need to know about this that is a little bit bizarre. If we take a look at a fully qualified domain name -- so that's my fully qualified domain name. Let's presume that client 1 has an IP address of 192.168.1.10/24, with a 24-bit subnet mask. We know that Client1 is the host name of this particular computer. Given my 24-bit subnet mask, I know that 10, just the fourth octet, represents that host on the network. For some reason with reverse lookup zones, they were like, "hmm, there's not a match". Over here, it's on the left. Over here, it's on the right. So here's what we'll do. We'll just write all the reverse lookup stuff backwards. A reverse lookup zone for this network would actually be written like that, with the three octets of the network ID being reversed. That's really the only trick to them. If you get into reverse lookup zones, look to make sure that it's written backwards.
Summary
1:22-1:27
Forward lookup zones, name to IP address. Reverse, IP to name, and they're always written backwards.

5.3.2 Standard DNS Zones

Standard DNS Zones
0:00-0:07
We're going to talk about standard zones. Standard zones are zones where the information is kept in text files on the DNS server.

Standard Primary Zones
0:08-0:21
There are two types of standard zones; standard primary zones have the only read-write copy of the zone. That means there's only one of them, and any changes that need to be made to that zone would be made on the primary.

Standard Secondary Zones
0:22-0:49
Standard secondary zones are read-only copies, and they're used for fault tolerance and to sort of balance the load. So if I'm using standard zones, I'm going to first create my standard primary. There's only one of them, and it's the only read-write copy of the zone. That's the most important thing to remember. That means that all changes to the zone are made on the primary. Secondary zones are read-only copies, and we're going to take a look at some of the rules on creating a secondary zone.

Zone Transfers
0:50-1:50
If you have a standard zone, the primary sends updates to the secondary using zone transfers, and there's a few things you need to know about zone transfers. In order to have a copy of the zone, it requires permission. And since changes can only be made on the primary, that's where I would go to set that permission. You have to be careful who you allow to have a copy of your zone.
Generally speaking, we name servers and clients with descriptive names so we know what their function is. If I can get a copy of your DNS zone, I have the names and IP addresses of all the computers in that domain, and probably I can figure out by looking which computers do what. So now as a hacker, not only do I have names and IP addresses, but I know who to target. So we're only going to allow servers that have permission to have a copy of the zone. Zone transfers are always initiated by the secondary, and it's based on the serial number on the SOA record, so it's always a pole operation.

SOA Record
1:51-2:49
The SOA record is the first record created in any zone. It identifies the zone. It identifies the primary server for that zone and a whole bunch of other things. On that SOA record is a serial number. It starts at one and goes to infinity. Anytime there's a change, the serial number is incremented. So when the secondary gets its first copy of the zone, it's got the same number in its SOA record as the primary. Every once in awhile, it will contact the primary and say "hey, I need your SOA record". If the number is different, it will initiate a zone transfer. The key is that it's initiated by the secondary, so if you need to run any commands to manually force a zone transfer, you're going to have to run those commands at the server that needs to get the changes. Because the zones are kept in text files, zone transfers are done in clear text. That means if I can capture a copy of your zone transfer, I can see all the information. If you have to have a standard zone and you have to secure it, you would use IPSec to secure the zone transfer.

Triggers
2:50-2:57
There are two triggers for zone transfers that will kick them off; one is the refresh interval. The other is called DNS notify.

-Refresh Interval
2:58-3:26
The refresh interval is just a period of time that goes by. By default, it's set to 15 minutes, so every 15 minutes, the secondary contacts the primary and says, "give me your SOA record". The primary sends that SOA record back to the secondary. If the serial number has changed, the secondary initiates a zone transfer and gets a copy of any changes that have occurred. If the serial number hasn't changed, it doesn't initiate a zone transfer.

-DNS Notify
3:27-3:48
DNS notify is a little bit different. In DNS notify, there's a notify list at the primary. If any changes occur, the primary will notify the secondary servers that are on that list. Bizarrely enough, the secondary server still says. "hey give me your SOA record". SOA record comes back, and it initiates a zone transfer.

Summary

3:49-4:18

So, standard zones are zones where the DNS information is kept on text files on the DNS server. There's two types; standard primary, which is the only read-write copy of the zone where all the changes occur, standard secondary, which is a read-only for fault tolerance and load balancing. Information gets back and forth between the primary and the secondary using zone transfers. Zone transfers require permission. They're always initiated by the secondary. They're clear text, so they're not secure, and there are two triggers: either the refresh interval or the DNS notify list on the primary.

5.3.3 Active Directory Integrated Zones

Active Directory Integrated Zones
0:00-0:20
Now, let's talk about Active Directory Integrated Zones. Active Directory Integrated Zones -- or some people say, ADI -- are zones that are actually integrated with Active Directory. That was sort of the sarcastic answer. That means that they're actually stored in the Active Directory database, that NTDS. DIT that lives on the domain controllers.

ADI Zones are Stored in Active Directory
0:21-0:28
Since Active Directory Integrated Zones are stored in Active Directory, that's the biggest thing to remember. If you could only remember that, you'd have a great starting place.

DNS Servers Must be Domain Controllers
0:29-0:41
That means that your DNS servers must be domain controllers. If the DNS server is not a domain controller, it can't be an Active Directory Integrated Zone, because something that's not a domain controller isn't going to get a copy of Active Directory.

Active Directory Replication
0:42-1:24
The DNS records are exchanged using Active Directory replication. There's no zone transfers at all when we're talking about ADI Zones. They go with replication. If we want to trigger "zone transfer," we would force replication. Active Directory replication is still a pole operation, but I would still do this at the domain controller that needs to get the updates, but, I would do it with Active Directory replication.
Active Directory is a multi-master, loosely consistent database. The key here is multi-master. When I make a change to Active Directory, I can connect up to any domain controller to make that change, and it will be replicated to all the other domain controllers in the domain. So at any given time, it's loosely consistent. Maybe not everybody's up to speed, but they're pretty close.

ADI Benifits
1:25-1:29
That's going to carry over into my Active Directory Integrated Zones.

Multi-Master
1:30-1:43
So Active Directory Integrated Zones, just like Active Directory, are multi-master. All the copies of the zones are read/write. So any situation where clients have to update records at more than one location, you've got to have an Active Directory Integrated Zone.

Secure Zone Transfers
1:44-2:02
Because the zone transfers go along with replication, they are secure, because Active Directory replication is encrypted by default. So if you're looking for secure zone transfers, you want to go Active Directory Integrated. We could do it with a standard zone, but we'd need to implement IPSec, and that's not exactly an easy thing to do.

Secure Dynamic Updates
2:03-3:48
They also are the only types of zones that support Secure Dynamic Updates. Secure Dynamic Updates mean only members of the domain can update records. Dynamic Updates allow DNS clients to contact the DNS server to update their records when their IP addresses change. You have to have Dynamic Updates supported in order to support Active Directory. There's just no way that a domain administrator could run around and keep DNS up to speed.
With Secure Dynamic Updates, whichever client creates the record becomes the owner of it. So client1 checks in with DNS, creates a record in the northsim.com zone. Client1 is the owner of that record. DNS can use the security id. of the computer to secure the record and limit updates to just client1. I can only do that with an Active Directory Integrated Zone, because the way I secure these records is using the information about that computer from Active Directory. Effectively what I can do then is limit who can update the zone to clients that are members of my domain.
Why is this important? Well, let's suppose I have a server, www.northsim.com, and it sells widgets, and everybody in the world is on there all day long. Kind of like an Ebay or a Best Buy, something like that. Some hacker decides, a great way to retire would be to get into DNS, change the records for www.northsim.com to the hacker's IP address, and temporarily reroute traffic to that hacker's imitation website, you know, probably able to get away with it for quite a while, and can collect a bunch of credit card numbers, bank information, before that person gets caught. So DNS securing it is very important, and having this ability to use Secure Dynamic Updates can really protect your network.

ADI Replication Scopes

3:49-4:58

The last thing that you really need to know about Active Directory Integrated Zones are the Replication Scopes. We can actually set up which domain controllers will have a copy of this zone. We've got four choices. One choice is All the Domain Controllers in the Domain, regardless of whether they run DNS or not. That was the way these worked in Windows Server 2000, and this choice is still in there for backwards compatibility.

A little bit better choice would be all the domain controllers that have DNS installed in the domain. If I have a situation where I need to provide name resolution for that domain across the forest, my third choice is all the domain controllers that have DNS installed everywhere throughout my forest. That means every domain controller that runs DNS will get a copy of this zone, and can authoritatively answer queries.

My last choice is an Application Partition. The only reason you would use this is if you want that domain information to be kept on less than all of the domain controllers in the domain. All of these choices, the smallest I can get is all the domain controllers in the domain; all the domain controllers with DNS. Here if I just wanted three out of six, five out of ten, I'd have to create my own application partition.

Partitions

4:59-5:38

The Active Directory database is divided into partitions. When you change the replication scope, what you're really doing is changing which partition you're going to store that DNS information in. The domain partition is already set up to go to all the domain controllers. That's my backwards compatibility option. There's a DNS domain partition that goes to all the DNS servers in the domain, and there's a forest DNS partition that's automatically already replicated to all the DNS servers in the forest. If I need to get down to less than everybody in the domain, I can create my own application partition inside that database, set up which domain controllers will receive a copy, and then store the DNS information there.

Directory Partition

5:39-6:04

To create my Directory partition, it's very simple. dnscmd /createdirectorypartition. Once I've got my partition created, I have to identify which domain controllers will get a copy. That's done using my dnscmd /enlistdirectorypartition, and then after that, I would specify who's getting a copy. I also could create the partition using my Active Directory Maintenance Tools; but if you're dealing with DNS, you probably want to stick with DNS command.

Summary

6:05-6:43

So Active Directory Integrated Zones are awesome because they're stored in Active Directory. They have to be on domain controllers. That's really the only sort of limiting requirement. They've got three benefits. They're multi-master. They're all primary zones. I can make changes on any copy. They have secure zone transfers because they're done with replication, which is already secure. And they can support Secure Dynamic Updates, which lets me limit who can update the record to just computers that are members of the domain. I've got four replications scopes, all the domain controllers in the domain, all the domain controllers with DNS in the domain, all the domain controllers with DNS in the forest, or if I want less than everybody in the domain, I'll make my own application partition.

Creating a New Zone
0:00-0:15
In this video, we're going to take a look at creating zones. The first thing I want to do is go into the DNS Management Console. We'll go up to Tools; DNS. The first decision that I need to make is, is it going to be a Forward Lookup Zone or a Reverse Lookup Zone?

Reverse Lookup Zones
0:16-1:14
We'll look at Reverse Lookup Zones first.
Reverse Lookup Zones map IP addresses to names. Forward does names to IP. I'm going to right-click and make a New Zone. The choices that you're going to see in this wizard are identical as the choices that I have for Forward Lookup Zones. We just want to see the one little thing that's a little bit different with Reverse Lookup Zones, and then we'll spend the majority of our time over on Forward Lookup.
I'm just going to say, Next, not worrying about zone type. Now I have to choos, is it IPv4 or IPv6, because I'm starting with an IP address, and I want to get back a name. I'm just going to leave it IPv4. Now, I put in my network ID. Let's say this is a Reverse Lookup for the 192.168.2.0 network. I'll hit Next. It's going to create a file for this. Fantastic. Next.
Reverse Lookup Zones: the main things you need to know about them are, they resolve IP addresses to names, and they're written backwards. If you know that, you should be in good shape.

Forward Lookup Zones
1:15-1:43
Now, we're going to take a look at Forward Lookup Zones. I'm going to make a new zone, and now we do want to go ahead and talk about this dialogue box. Right now, this particular server is a member server. Member servers can only host standard zones. If I want to do Active Directory Integrated, which is this checkbox down here, I've got to be on a domain controller.
When I'm creating a standard zone, I have to decide, is it the Primary zone, Secondary zone, or a Stub zone?

Primary, Secondary and Stub Zones
1:44-2:27
Primary zone is the only read/write copy of the zone when you're dealing with standard zones. Secondary zones or Stub zones are read-only. Stub zone is a copy of the zone that only contains DNS server information. It's used a little bit differently than Primary or Secondary. If you're just looking at name resolution for the zone, you're really just looking at Primary or Secondary.
We'll do a Primary since this is going to be the first copy. I'll hit Next. Now I provide the name of my zone. Next. Since this is a standard zone, the DNS information is kept in a file on the hard drive of the server, so it's telling me the name of the file. It's just the name of the zone.dns, and that's fine. I don't need to mess with that.

Updates
2:28-2:59
Now, it asks me if I want to allow updates, and you can see that by default, it's saying do not allow updates. My only choice is to allow both nonsecure and secure. Notice that secure dynamic updates is grayed out. That's a feature that requires Active Directory Integrated Zones. I don't have an Active Directory Integrated Zone because I'm not a domain controller, so I can't enforce that. The best I can do is to say I'll take nonsecure updates and be happy with that. Now, I hit Finish, and I've got my new zone.

Standard Zone
3:00-3:04
With standard zones, the Secondary zone has a copy of the zone.

Zone Transfers
3:05-3:14
When you're dealing with a standard zone, the information gets from the Primary to the Secondary using zone transfers, and there's three things you need to keep in mind about zone transfers.

Permission
3:15-4:37
The first thing is, they require permission. If I come in here, I've got to make sure there's a checkmark in Allow zone transfers. I can allow them to servers listed on the Name Servers tab, which is up here. Or, more secure, I could provide the IP address or FQDN of the server that's allowed to have a copy of the zone. Either way, I need to provide permission before I create the Secondary. That permission can only be configured on the Primary, because the Primary is the only read/write copy.
If I actually want another server to be allowed to have a copy of this, I've got to go into Name Servers, Add in another server, and I'm getting an error here because I don't have a copy of the zone over on dc1.northsim.com. It's kind of the chicken or egg type of thing. You should give permission before you create the zone, but it's yelling at me because there's no zone over there. If I created the zone first, the creation of the zone will fail, and I'll get an error until I give permission.
In practice, it's better to give permission first, because what I find is that if I don't give permission first, I end up having to delete the zone and recreate it, and it's really annoying. Even though it's yelling at me, that's okay. It's not a problem. I'm just going to go ahead and Apply. That gives that computer permission to do a zone transfer, and it needs to do a zone transfer to get a copy of the zone.
Triggers
4:38-5:05
Zone transfers are triggered by one of two things: One is the refresh interval. That specifies how many minutes go by before the Secondary will check in with the Primary to see if there have been changes. Changes to the zone are tracked using the serial number, and that's how the Secondary knows if there's been a change to the zone, and it initiates another zone transfer, if I don't want to wait.
I have a particular server with the Secondary zone. It should always be up to date.
DNS Notify List
5:06-6:36
The other way to do this is using the DNS Notify list, which is inside this Notify button. In here, I would list out the DNS servers that have a copy of the zone that should be proactively notified immediately if there are any changes. I'm just going to go ahead and click OK.
Now that I've given permission to DC1, I can go over to DC1 and actually create the zone. This time, I'm creating a Secondary zone. You notice, as soon as I click on Secondary Zone, it grays out Store the zone in Active Directory. Active Directory Integrated Zones are all Primary, because Active Directory is a multi master data base. There's no such thing as an Active Directory Integrated Secondary zone, because they are multi master. I'm going to hit Next, and now, I've picked up a copy of that zone.
If I hadn't given permission in advance, this would have been read, and then you've got to try to transfer from Master Reload. Half the time, it doesn't work, so better to give permission first.
Active Directory Integrated Zones can only be created on domain controllers, and the information is actually stored in Activity Directory. I will leave this checked. It's got to be a Primary zone, or a Stub zone would be fine, but that just has the DNS information.
If I do Active Directory Integrated, I have four scopes that I can use for replication.
Scopes
6:37-6:58
To all the domain controllers in this domain, whether they have DNS or not, which you can see is used for Windows 2000 compatibility, the default is to all the DNS servers running on domain controllers in this domain, but I could also replicate it to all the DNS servers running on domain controllers in this forest.
Directory Partition
6:59-7:13
If I want less than all of the domain controllers, I would need to create a directory partition, and then I could store the zone on that partition. We're just going to take the default, and hit Next. Now, that I get to my Dynamic Updates page, notice I can select "Allow only secure dynamic updates" because that's available for Active Directory Zones.
Dynamic Updates
7:14-7:22

Summary

7:23-8:00

My three selling points for Active Directory Zones are multi master. I can update records at any copy of the zone. Secure dynamic updates, which means that updates can be limited to members of the domain, and secure zone transfers. Standard zones are kept in text files, and the zone transfers go through in clear text. Here, because the information is stored in Active Directory itself, and Active Directory replication is automatically encrypted, the DNS information as it goes along with Active Directory, of course, is going to be encrypted as well. Now I have a new Active Directory Integrated Zone.

That's how we create standard zones and Active Directory Integrated Zones if we need to create them manually.

5.3.5 Configuring Zone Properties

Configuring Zone Properties
0:00-0:36
In this video, we're going to go through the properties of DNS zones. I'm going to go into my DNS console snap-in, and we can pretty much pick any zone here to take a look at.
Now, because this is a secondary zone, it doesn't have all of the information that the primary would. You can see that the SOA record is grayed out. Standard zones can only be modified on the primary zone. Make sure you're on a primary zone or an Active Directory-Integrated if you want to configure the properties of the zone.

General Tab
0:37-1:14
The General tab is used to switch the type of zone. You can see that currently I have an Active Directory-Integrated zone. If I wanted to change that, I can get back into that first page of the wizard that we saw in creating the zone. If I don't want it to be Active Directory-Integrated, I could uncheck it. I could convert it into a secondary zone or I could convert into a stub zone.
If it is Active Directory-Integrated, I can change the replication scope and you can see your four replication scopes listed here on this page. The fourth one is grayed out because I haven't created a directory partition. If I had a directory partition, I could store it on that partition.

Dynamic Updates
1:15-1:41
Here I have my settings for Dynamic updates. If you want the computers to be able to update their own records in the zone, you need to have it either on Secure only or Non-secure and secure. If you have it set to None, that's going to create a problem, particularly for Active directory. Active Directory requires that the clients can update their own IP addresses. The Secure only choice is going to be available for Active Directory-Integrated zones. If it's a standard zone, I won't be able to do that.

Aging and Scavenging
1:42-1:59
Aging and scavenging is used to automatically remove old or stale records. I can go through and set it up to scavenge stale records. The No-fresh interval means the time between the most recent refresh and the moment when it may be refreshed again.

Refresh Interval
2:00-2:21
Refresh interval is the time between when it's refreshed and when the record can actually be deleted.
The important thing to know about aging and scavenging is that it has to be turned on, on both the zone and the server. Even if I turn it on here, the stale records aren't going to get removed from that zone. I've got to enable aging and scavenging on the server before anything's going to happen.
The Start of Authority record is actually built into the Properties of the zone.

Start of Authority Record
2:22-3:26
It's actually a record in the zone, but it shows up in the properties. We've got the serial number, which gets incremented every time there's a change at the zone. Notice it has the email address of the responsible person. Refresh interval is how often we're going to try to have a zone transfer. "Retry interval" is, if that zone transfer fails, how long goes by before the secondary is going to try again. These are not really in play with Active Directory-Integrated, because those go along with Active Directory replication. These are more for standard zones.
"Expires after:, is how long I can go without a zone transfer. What we're looking at now tells us that if we go one day without any type of an update from the primary, then this DNS server would stop responding to queries. This is how long you need the DNS server to work without having a zone transfer.
Time to live is the time to live that would be placed in the DNS cash on any records that are resolved from this zone. This is the time to live for the SOA record itself.

Name Servers
3:27-3:33
The Name Servers tab lists all the DNS servers that have a copy of this zone.

Zone Transfer

3:34-4:01

Zone transfers is used to set up who's allowed to have a copy of the zone. If I don't permit zone transfers, it means nobody can have a copy of the zone. This needs to be done before you set up the secondary. It only applies to standard zones. You can see here, I'm on Active Directory-Integrated zone. It's unchecked. It's irrelevant. Who is going to get a copy of this zone is going to go by my replication scope. Zone transfers are only in play for standard zones.

WINS

4:02-4:17

WINS allows the computer to actually forward queries to a WIN server for netBIOS name resolution. This is very old stuff. Microsoft would prefer that you make a global name zone and call it good. We're not going to worry too much about this tab.

Security

4:18-4:39

Security is used to allow people to create and manage records in the zone. I say people, because by default, the computers have enough rights to come in and update the zone. You would really just use this if you have a specific user that you want to be able to make records in the zone. In general, we don't do that. We set up DNS, you walk away, it just runs and runs and runs no problem.

Summary

4:40-4:52

Those are all the different properties of the zone. Make sure you're particularly familiar with the information on the SOA record in the General tab, so that you know what type of settings can be set up on the zone.

5.3.6 Zone Management Facts

An authoritative DNS zone holds a full copy of the DNS records for a zone. The table below lists the types of authoritative DNS zones:

Zone Type	Description
Primary	The primary zone contains the master copy of a zone database. • The primary zone is the only writeable copy of the zone database. • Changes to the zone can be made only to the primary zone database. • The server that holds the primary zone is called a *primary server*. • Each zone can have only a single primary zone server. • Zone data is stored in a text file. • The primary zone sends updates to the secondary zone using zone transfers. • Permission is required to have a copy of the zone. The permission should be set in the primary zone.
Secondary	A secondary zone is a read-only copy of the zone database. • Changes cannot be made to the records in a secondary zone. • A server that holds a secondary zone is called a *secondary server*. • Secondary servers receive copies of zone data from other servers using zone transfer. • Secondary servers can receive zone data from the primary server or other secondary servers. • Zone transfers are always initiated by the secondary zone. The zone transfer is based on the serial number of the SOA record. The SOA record is the first record created for a zone. The SOA record identifies the zone and the primary server for the zone. The serial number for the SOA record is incremented whenever there is a change to the SOA record. • Triggers for zone transfers are: **Refresh interval** specifies the amount of time between requests for the SOA record of the primary zone. **DNS Notify** lists the servers to be notified. The primary server sends a notification to the secondary server that a change has been made. The secondary server then initiates a zone transfer by requesting a copy of the SOA record. • Zone data is stored as clear text. You can use IPSec to secure the zone transfer.

Active Directory-integrated	An Active Directory-integrated (ADI) zone holds zone data in Active Directory instead of a text file. • Active Directory-integrated zones are multi-master zones, meaning that changes to the zone information can be made by multiple servers. Multiple servers hold read-write copies of the zone data. Only DNS servers that are domain controllers can host Active Directory-integrated zones. Zone data is stored in Active Directory. Replication of zone data occurs during Active Directory replication. Storing zone data in Active Directory provides automatic replication, fault tolerance, and distributed administration of DNS data. Zone transfer is secure, because Active Directory replication is encrypted. • Active Directory-integrated zones support secure dynamic updates. Dynamic updates allow DNS clients to contact the server and update their records when their IP address changes. Only members of the domain can update records. The client who created the DNS record becomes the owner of it. • You can configure a secondary server to get zone data from an Active Directory-integrated zone. However, you cannot have a primary zone and an Active Directory-integrated zone for the same zone. • Active Directory-integrated replication scopes allow you to specify the domain controllers that will have a copy of the zone data. The choices are: All domain controllers in the domain, even if they are not running DNS. All domain controllers in the domain that have DNS installed. All domain controllers in the forest that have DNS installed. An application partition that allows you to choose the domain controllers that will have a copy of the zone data. ▪ You use this to specify which domain controllers will have zone data. ▪ The Active Directory database is stored in partitions. When you use application partition, you change the partition that the zone information is stored in. ▪ To set up an application partition: ▪ **Dnscmd /createdirectory** creates the partition. ▪ **Dnscmd /enlistdirectorypartition** specifies the domain controllers included in the application partition.

The zone types above describe the read-write capabilities and the storage location of zone data. In addition, zones are classified as one of two types:

- A *forward lookup zone* provides hostname-to-IP address resolution. Clients query the DNS server with the hostname and receive the IP address in return.

- A *reverse lookup zone* provides IP address-to-hostname resolution. Clients query the DNS server with the IP address and receive the hostname in return. In a reverse lookup zone, the octets in the zone name are written in reverse order. For example the reverse lookup zone for the IP address of 192.168.1.10 / 24 will be written 1.168.192.in-addr.arpa.

5.4 DNS Records

As you study this section, answer the following questions:

- What is the difference between an A and a quad-A record?
- What situations might warrant using a CNAME record?
- What type of server does an MX record identify?
- What is the purpose of PTR records?
- What does an NS record identify?

After finishing this section, you should be able to complete the following tasks:

- Create A records and associated PTR records for specified hosts.
- Create a CNAME record to be used as an alias to redirect requests.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  Create DNS Records
  - Create a Zone and Add Records
  - Create A and CNAME Records

This section covers the following 70-410 exam objectives:

- 403 Deploy and configure DNS service.
  This objective may include but is not limited to:
  Create A and PTR resource records
- 501 Install domain controllers.
  This objective may include but is not limited to:
  Resolve DNS SRV record registration issues

5.4.1 DNS Record Types

DNS Record Types
0:00-0:06
Let's talk about some common DNS Record Types. You want to have a basic idea, very simple, what each type of record does?

A Records
0:07-0:13
A Records take a fully qualified domain name and map it to an IPv4 address. Make sure you know it's an IPv4; if it's IPv6, then you need a AAAA Record, and the easiest way to remember this is an IPv4 address is 32-bits.

AAAA Records
0:14-0:28
IPv6 is 128-bits -- well, that's four times as big, and so is the DNS record.

PTR Records
0:29-0:49
PTR Records are pronounced Pointer Records, and they map IP addresses to names. Make sure you know that it's either type of IP address, so there's not a separate record for IPv6. Either or, both would be a pointer record. A and AAAA Records go in forward lookup zones. Pointer Records are the types of records that we create in reverse lookup zones.

NS Records
0:50-1:02
NS Records identify DNS servers or the zone; so an NS record says, okay, this is the server who is a DNS server for that particular zone and there'll be one NS record in the zone for each DNS server that has a copy of that zone.

MX Records
1:03-1:15
MX Records identify email servers. Now, these records are a little weird, because they have a priority. Lower priorities are the preferred server, so a server with a 10 Priority is going to get the traffic over a server with 20 Priority.

SPF Records
1:16-1:22
SPF Records identify authorized email servers, and they're created using TXT Records.

Email Spam
1:23-2:49
One of the biggest problems with email is spam, so initially, before we had such spam issues, all we needed was an MX record to figure out who the server was for that email, so if I wanted to send somebody email -- shad@northsim.com -- I just need an MX Record that says mail.northsim.com is the email server for that domain. Well of course, spammers began installing email services and telling the email service, "yeah, I'm the mail server for northsim.com", and sending out email using that identity. Initially what they would do is reverse lookup, so if the email server claims its name is mail.northsim.com they could go to the pointer record for mail.northsim.com, and make sure that the IP address of the server that was sending the email was really the same as that post name. The problem comes in when they started to go ahead and put in fake mail names, so the spammers really do name the server mail.northsim.com and managed to get a reverse lookup record put in there. How do I know, then, that this really is the authorized server for that domain? That's where the SPF Record comes in. I can use the SPF record to find out which host on the internet is allowed to do that when I have a receiving email coming in, and then make sure that that's the correct email server. MX records are really more used for outgoing mail. They let me know who to send it to, if I'm sending email to that domain. SPF records are used to verify the authorized email servers for incoming email to make sure that they are not spammers.

CNAME Records

2:50-3:54

CNAME records are for aliases. Let me give you a little example suppose I want to go ahead and create a website, www.northsim.com, but Northsim is a small company, and they don't want to host their own website, so they go to their ISP and they hire them to host the website. Let's just say that the server that the ISP puts it on is named web1, so this server web1.isp.com -- if it changes its IP address, it will update the A Record in the isp.com domain, but we need anybody who's going to www.northsim.com to be directed to web1.isp.com. It would be possible just to go in and make an A record and use the IP address of the web server, but if the web server's IP changes, I'll have to manually update that record. A CNAME record would look something like this: it would be in the northsim.com, domain and it would say, if you're looking for www, go see the name web1.isp.com. That way, no matter how many times web1 changes its IP address, anybody headed for www.northsim.com is still going to be directed to the right spot.

SOA Recrods

3:55-4:18

SOA records are Start of Authority records, and they have a bunch of things on there. It's the first record created in the zone. There's the Serial Number, which is used to track changes in the zone. It also identifies the responsible administrator by email address, so for every zone there should be an email address in there. This is who you email if you have an issue with DNS for that particular domain. It also has all the settings for the zone on the SOA Record as well.

Summary

4:19-4:26

Those are some of the most common DNS Record Types, and you just want to be familiar with them so that you know which type of record to create for whatever your situation is.

Creating Common Records
0:00-0:18
In this video we're going to take a look at creating some common types of DNS records. To do that, we need to open up the DNS console. First of all, we can take a look. In a forward lookup zone, A records map names to IPv4 addresses.

A Record
0:19-0:54
If I needed to do a new A record, I just right click and do a new host. If I do have a reverse lookup zone, I can have it created; an associated pointer record, which maps the IP address to the name. Then, I can also set up who's allowed to update DNS record.
I'm going to go ahead and add this host. If I put in an IPv6 address, it would be a quad A record. It's the same box, but you should know that quad A (AAAA) is named IPv6; simple A record is named IPv4. Those are important.

MX Record
0:55-1:31
Another common type of record would be an MX record. MX records identify email serves. If I don't specify any domain up here, I'm talking about the email server for northsim.com, and then I would specify who that is. Notice MX records have a priority. The lower the number, the higher the priority of the server. If there were two MX records for northsim.com, 1 out of 10, 1 out of 20, all the email traffic would go to the one with the 10, unless that server was down, and then it would go to the one with the 20.
Another very common type of record is the CNAME record, which is used for aliases.

CNAME Record
1:32-3:04
If you take a look, Member2 is actually the server that has the IP address 192.168.2.51. I created an A record for www to point to the same IP address. That computer can't have two names. It's either named Member2, or it's named www. In this case, it's actually named Member2. If I change the IP address of Member2, it's not going to update the www record.
If I'm hosting a website on that computer, this is a bad way to do it, because every time Member2's IP address changes, I'm going to have to come into DNS and manually update this record. If you have a computer that should be known by multiple names, you do one A record in the real name of the computer, and then the other names you use the CNAME record for.
I'm going to delete my erroneous host record here and correctly configure a CNAME record. I would say, look, if you're looking for the computer www.northsim.com, go look for member2.northsim.com. Now it's perfect. Anybody going for www.northsim.com is going to get directed to the member2northsim.com A record, which is here, and that's the record that the computer is going to keep up to date whenever it's IP address changes.
You can also see, in here, we have NS records that identify who the DNS server is for this particular domain. This domain has one DNS server, and the domain is hosted on dc1northsim.com.

Start of Authority (SOA) Record
3:05-3:42
The Start of Authority (SOA) record, which also shows up in the properties of the zone, is the first record in the zone, and it has all the zone information. I can either go to the Properties of the zone, or if I double click this, it takes me right into the Properties.
There's the Serial number, which gets incremented when there's changes; the name of the Primary server that's hosting the zone, email address of the Responsible person who's in charge of the zone, and then my zone settings.
Those are some of the common types of records that you'll run into; not difficult to make, but definitely important to have an idea of what each type of record is used for.

5.4.3 DNS Record Facts

Entries for hostnames, IP addresses, and other information in the zone database are stored in *records*. Each host has at least one record in the DNS database that maps the hostname to the IP address. The following table lists common resource records.

Record Type	Use
SOA (Start of Authority)	The first record in any DNS database file is the SOA. The SOA record: • Defines the general parameters for the DNS zone. • Is assigned to the DNS server hosting the primary copy of a zone. • Is the first record in the zone database file. There is only one SOA record in each database. • Includes parameters such as the authoritative server and the zone file serial number. • Includes an email address for the administrator responsible for the DNS domain.
NS (name server)	The NS resource record identifies all name servers that can perform name resolution for the zone. Typically, there is an entry for the primary server and all secondary servers for the zone (all authoritative DNS servers).
A (host address)	The A record maps an IPv4 (32-bit) DNS host name to an IP address. This is the most common resource record type.
AAAA (quad-A)	The AAAA record maps an IPv6 (128-bit) DNS host name to an IP address.
PTR (pointer)	In a reverse lookup zone, the PTR record maps an IP address to a host name by pointing to the appropriate "A" or "AAAA" record. • IPv4 PTR records are created in the in-addr.arpa namespace. • IPv6 addresses are created in the ip6.arpa namespace.
CNAME (canonical name)	The CNAME record provides alternate names (or aliases) to hosts that already have an A record. This record enables a server to be referred to by different names in DNS. These records are useful when: • Migrating servers. • Referring to a server a complex name with a user-friendly name. • Redirecting traffic to an ISP hosting the site.

MX (Mail Exchanger)	The MX record identifies servers that are available to receive Simple Mail Transfer Protocol (SMTP) mail. Low priority indicates the preferred server.
SRV (service locator)	The SRV record is used to indicate the resources that perform a particular service. This allows clients to find services (such as domain controllers) through DNS. Windows automatically creates these records as needed.
WINS and WINS-R resource records	Add these records to a zone when you want to allow DNS to use WINS resolution. The WINS resource records identify the WINS servers to forward failed DNS resolve requests. The WINS-R resource record allows the resolution of a reverse query that is not resolvable through DNS.
SPF (Sender Policy Framework)	SPF records identify authorized email servers. SPF records are created using TXT records. DNS uses the SPF record to verify that the host sending the mail is authorized to use the DNS name.

5.5 DNS Server Properties

As you study this section, answer the following questions:

- When should you activate debug logging?
- Where would you enable scavenging of stale records on the system?
- What is the difference between a simple query and a recursive query against the DNS server?
- What type of DNS servers are allowed when the *Enable BIND secondaries* option is checked?
- Which advanced DNS option is used to verify that the DNS record was not modified in transit?

After finishing this section, you should be able to complete the following task:

- Configure DNS server properties in DNS Manager.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  Manage DNS Configuration

This section covers the following 70-410 exam objective:

- 403 Deploy and configure DNS service.

5.5.1 DNS Server Properties

DNS Server Properties
0:00-0:05
We are going to talk a little bit about the properties of the DNS server. We want to talk about the settings on the Advanced tab.

Disable Recursion
0:06-0:21
First option we're going to see on there is the option to Disable Recursion, which also disables forwarders. This means that the DNS server cannot make recursive queries, it can't be a DNS client, it can't do forwarding, and you would just enable it for security if you know that the server doesn't need those functions.

Enable BIND Secondaries
0:22-0:46
Another option on there is Enable BIND secondaries. Whenever you see the word BIND like that, all in uppercase, that means Unix DNS. It's possible that you might need to set up your Microsoft DNS server to interact with the Unix DNS server. Certain versions of Unix were a little bit slower than the newer versions. What BIND Secondaries does is allows your Microsoft server to interact with the older, slower version.

Secure Cache Against Pollution
0:47-1:36
'Secure cache against pollution' does exactly that. If you've ever gone to a website and when you look up at the address tab in Internet Explorer or whatever browser you're using, and you see a different domain name, that just means you've been redirected to that new domain name. I tried to go out to www.northsim.com and look up, and it's eastsim.com. The concern is, northsim.com has been hijacked, and that's why I was redirected. Secure cache against pollution, which is enabled by default, just says this: if the DNS server sends out a query for one domain and gets back a response for a different domain, it will not put that information in the DNS cache, just in case the first domain has been hijacked. That way, that information won't stay in the cache, and as soon as the problem is corrected, that server would re-resolve the name and not have any issues. Enable DNSSEC validation is also in there.

Enable DNSSEC Validation
1:37-2:12
This is new with Windows server 2012. DNSSEC is a technology that allows the DNS server to sign the DNS records. The clients can validate that signature and make sure that the record wasn't modified in transit. I can't sit outside your company with a packet sniffer and replace the DNS queries for your website to buy things from your company with my IP address. This is only supported by Windows server 2008 R2 and Windows 7 clients by default. This setting lets me use DNSSEC Validation for DNS domains outside of my domain.

Netmask Ordering
2:13-3:16
These two are important. Netmask Ordering and Round Robin both deal with multiple records. I put out a request for www.northsim.com and there are multiple A records with IP addresses. The preference is towards Netmask Ordering. The client receives the record that matches its class of IP address. Let's take a look at that. Maybe my DNS server in the northsim.com domain has two records for the www server. One record says the IP address is 192.168.1.10. The other one says the IP address is 10.1.1.40. Why would I have two records? Maybe there's so much traffic for that server that I need to have a couple of different servers that can respond to queries. It's done for fault tolerance and load balancing. If the client comes in with an IP address that's class C address, so something else in the 192.168.1 network, it would receive this record. A client coming in with a class A address, like 10, will get that record, the idea being that I want to give back the record that's going to direct the client to the server that's closest to them.

Round Robin
3:17-3:34
If the DNS server can't match one of the records to the IP address of the clients, then it's going to go to round robin. All that means is DNS is going to rotate through the records. The first request gets record 1, second gets 2, third gets 1, fourth gets 2, and it just rotates through them.

Aging and Scavenging

3:35-4:19

Last thing we're going to see on that Advanced tab is Aging and Scavenging. Basically, what this does is, it will automatically remove old records. Sometimes they might say stale, they might say outdated ... however we talk about it, these are records that have not been updated by the clients in quite a long period of time. We can automatically remove them from the zone to keep the zone file small. Microsoft actually doesn't recommend turning this on unless you need it, but it's important to know, if you do want Aging and Scavenging to work, it must be set on both the server and the zone. On that Advanced tab, we can turn it on. We're still going to need to go into the properties of the zone and enable it there as well, in order for those records to actually be removed. Those are some of the important properties inside of the properties of the DNS server.

Configuring DNS Server Properties
0:00-0:20
In this video, we're going to look at the properties of the DNS server. I need to go into DNS Management Console, and we're going to go into the Properties. You want to keep your information here very streamlined, so we'll go through the different tabs. None of it is very difficult. It's just whether you've seen it or not before.

Interfaces
0:21-0:49
Interfaces specifies the IP addresses on which DNS will respond to DNS queries. If I had a particular IP address I didn't want it responding on, I could uncheck it. That does not lock down who can use this server for DNS. The only thing that can control who can talk to the server is a firewall. All this says is, I'm going to respond to queries that come in on this IP address. It doesn't say anything about where the computers that are sending the queries actually live.

Forwarders
0:50-0:55
Forwarders is used to forward traffic, for which I'm not authoritative. Root Hints lists all the root servers on the Internet.

Root Hints
0:56-1:01
Debug Logging is used when we want to do extra logging for troubleshooting, so anything that's out of the ordinary, and, you can see, we also have an Event Logging tab that says, hey, we're going to log Errors -- Errors and warnings.

Debug Logging
1:02-1:28
These are ordinary events. Any time you're looking for any advanced logging, you see we can specify particular packets; Ingoing, Outgoing. Then, any type of extra logging is Debug Logging.

Monitoring
1:29-1:57
The Monitoring tab lets me run simple query against the DNS server or recursive query. A simple query says, run a query against one of my zones. Make sure it works. Recursive query would test forwarding, and it's going to fail because I don't have forwarding set up, and you can see that the recursive query failed, which is fine. We just looked at the Forwarders tab, and there's nothing in there.
The Security tab is used to give people permission to modify the properties of the DNS server.

Security
1:58-2:03

Advanced Tab
2:04-2:10
Pretty much, the only tab we have left, which is the most important tab in here, is the Advanced tab.

Disable Recursion
2:11-2:27
Disable recursion disables Forwarding and disables the DNS client. It might gain you a marginal amount of security, so if you just know it's not going to do forwarding, it shouldn't be a DNS client. You can come in here and check this. You can see it's not checked by default. It's not something we desperately need to turn on.

BIND (UNIX DNS)
2:28-2:41
Any time you see the word BIND like that, all in upper case, it means UNIX DNS. If I check this, it's going to allow my Microsoft DNS server to communicate with certain UNIX DNS servers.

Fail on Load if Bad Zone Data
2:42-2:51
"Fail on load if bad zone data" is exactly what it sounds like. If the zone data is bad, don't load it up.
I'm going to drop down a couple. Secure cache against pollution, which is turned on by default, does exactly that.

Secure Cache Against Pollution

2:52-3:31

It tries to make sure that bad entries don't get into the DNS cache; if you've ever gone out on the Internet to one domain, and then it jumps to another one. You go to www.shad.com, and it jumps to www.shadow.com. If your DNS server makes a query for one domain and gets back an answer for a different domain, the concern is that the original domain has been highjacked. In that case, it doesn't put the answer in the DNS cache. That way, it won't stay in the cache and maintain that high jacking beyond the time when it's fixed. That's all it does.

Enable DNSSEC Validation for Remote Responses

3:32-4:04

This last checkbox Enables DNSSEC validation for remote responses. DNSSEC basically allows the DNS server to digitally sign the records that it sends out to the clients, and the clients can verify that signature to make sure that the DNS record was not modified in transit. That came in with Windows Server 2008 R2, I believe, so it's not going to be supported by all clients. It's supported by Windows 7 and Windows 8. Anything above that will support it. Anything below Windows 7 is not going to support DNSSEC.

Round Robin and Netmask Ordering

4:05-4:34

The last two left in this list are round robin and netmask ordering, and you can see they're both checked by default. netmask ordering says this: If there's multiple A records or multiple of the same type of record for the request, go ahead and give the client the record that most closely matches their IP address.

If you can't resolve it that way, then it will default to round robin, where it'll just cycle through the records. You don't have to worry about Name checking or Load zone data on startup.

Enable Automatic Scavenging of Stale Records

4:35-4:51

You should take a look at "Enable automatic scavenging of stale records". That's my aging and scavenging; it allows me to have DNS automatically remove records that haven't been updated within a certain amount of time. The key to aging and scavenging is it has to be set on both the server and the zone.

Let me just show you the multiple records real quick.

Multiple Records

4:52-5:52

You can see here, I've got an A record for dc1. There's nothing to prevent me from making another A record for dc1 as well. You might think, why would we see that? Well, maybe dc1 has two network adapters, and it's registering two addresses. The key is this: Because there's now multiple records for this server, when the request comes in, if it comes from a computer whose IP address starts with 192, that computer's going to get this record. If the computer's IP address starts with a 10, it would be given this record.

Let's suppose the client comes in with a 172 address. There's not a record that really matches that. In that case, we'll default to round robin, where it'll just; this one, this one, this one, this one. It just cycles through them, one, two, one, two, one two, one two. It's kind of a form of rudimentary load balancing that can be done with DNS.

Summary

5:53-5:59

Those are the properties of your server. Make sure you're familiar with those options on the Advanced tab and some of the other tabs that are in there, and you should be good to go.

5.5.3 DNS Server Properties Facts

Use DNS Manager to configure DNS server properties.

Tab	Description
Interfaces	The **Interfaces** tab identifies the IP addresses DNS will use to listen for DNS queries. Uncheck IP addresses that you do not want DNS to use for queries. Choosing this option does not restrict the use of the IP address. The only mechanism that can restrict traffic to the DNS server is a firewall.
Forwarders	The **Forwarders** tab identifies DNS servers that queries are sent to when the local DNS is not authoritative.
Root Hints	The **Root Hints** tab identifies the root servers on the Internet.
Debug Logging	The **Debug Logging** tab allows you to set detailed logging parameters for advanced troubleshooting, including: • Packet direction and transport protocol • Packet content and packet type • Other options: Log unmatched incoming response packets Details Filter packet by IP address • The log file location and maximum size
Event Logging	The **Event Logging** tab allows you to specify events to be logged. Options include: • No events • Errors only • Errors and warnings • All events
Monitoring	On the **Monitor** tab, you can run a simple or recursive query against the DNS server. A recursive query tests forwarding.
Security	On the **Securities** tab, set permissions for users or groups to modify the DNS server.

Advanced	Server options on the **Advanced** tab include: • **Disable recursion** disables forwarding and the local DNS client. • **Enable BIND secondaries** allows the Microsoft DNS server to communicate with UNIX DNS servers. • **Fail on load if bad zone data** prevents the zone data from loading if it is found to be incorrect or corrupt. • **Enable round robin** enables cycling through records when there are multiple records of the same type with the same name. • **Enable netmask ordering** provides the requester with the record that most closely matches the requester IP address when there are multiple records of the same type with the same name. • **Secure cache against pollution** protects the cache by not updating an entry when a request sent to one domain results in a response from another domain. • **Enable DNSSEC validation for remote responses.** DNSSEC allows DNS servers to digitally sign the records it sends out to the clients. The clients can use the signature to verify that the DNS record was not modified in transit. This option is available only on clients running Windows 7 and later. • **Enabling automatic scavenging of stale records** removes records that haven't been updated within a specified period of time. Aging and scavenging must be set on both the server and the zone. In the zone properties, the scavenging of stale records is enabled or disabled. In addition, the **no-refresh interval** and the **refresh interval** parameters can be configured. The no-refresh interval specifies a time period where updates to DNS records are not allowed. After this time period, the record can be updated for the period of time specified by the refresh interval. Scavenging only removes stale records that have not been updated after the refresh interval has expired.

5.6 DNS Troubleshooting

As you study this section, answer the following questions:

- What does it mean when a DNS server reports that a hostname could not be found?
- When troubleshooting with ping, how can you determine if the problem is with the DNS server or with your client?
- What tool can you use to view specific records on the DNS server?
- How do you clear the DNS cache?

After finishing this section, you should be able to complete the following tasks:

- Use Ping and NSLookup to troubleshoot DNS.
- Use IPConfig to manage the contents of the DNS cache.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 4.0 DNS.
  Create DNS Records
  - Troubleshoot Name Resolution

This section covers the following 70-410 exam objective:

- 403 Deploy and configure DNS service.
  This objective may include but is not limited to:
  Manage DNS cache

DNS Troubleshooting
0:00-0:08
Let's talk about DNS troubleshooting. I do most of my DNS troubleshooting from the command line, and I've got a few commands that you can use to do that. Let's take a look at some of them.

Ping
0:09-0:12
The first command that I use to troubleshoot DNS is Ping.

Ping by Fully Qualified Domain Name
0:13-0:58
If you ping a fully qualified domain name and what comes back is a line of text essentially saying I couldn't find it, that right there tells you it's DNS. If it was a problem with TCP/IP, you'd either get no response or destination host unreachable, which indicate different problems. So right away, when I get back a line of text, I know DNS is not working.
My next question is, is it a problem with my client contacting the DNS server, or is it a problem with the DNS server itself? There's a couple things that I can do; one thing I can do is to ping my DNS server by IP address. If my DNS server will answer me if I ping it by IP address but I'm not getting an answer on that client, then the problem is somewhere on DNS.

Ping by IP Address
0:59-1:24
The other thing that I could do is go to another computer, find out the actual IP address of the fully qualified domain name I'm trying to contact, and ping that. Again, anytime you can ping by IP address and not by name, it's got to be a problem with DNS. The fact that I can ping by IP address tells me everything's good with TCP/IP between me and that other host; I'm just not getting name resolution.
Another command that you can use is NSLookup.

NSLookup
1:25-2:18
NSLookup is a command for querying DNS. It's really complex. For example, you can go in and specify which DNS server it should send the query to. You can tell it what types of records you'd like to look at. You can really define that query as much as you need to. Be aware, of course, that some queries require permission. So if I were to contact let's say Microsoft, and say, "give me a copy of all your records", well, that's almost the same as a zone transfer, so I'd need permission to be able to do that. But NSLookup is a very detailed command that's going to get me responses directly back from the DNS server. So if I've isolated this down to a problem with the DNS server and I find that the DNS service on the server is running but it's just this particular record that's a problem, I would use NSLookup to query and see what the DNS server has as far as those records go.

IPConfig
2:19-2:58
The last command that I use, which is really helpful would be IPConfig, and there's three switches that will help you with DNS. /displaydns will show you the contents of the DNS cache. /flushdns clears out that DNS cache. So if you think the problem is that the host IP address has changed and it's stuck in the cache, you can do a /flushdns to clear that cache out. If it won't leave the cache, make sure you check the host file. If the problem is that this particular client's IP address is wrong in DNS, ipconfig /registerdns causes the client to contact DNS and re-register its IP address. So those are the three main tools that I use for troubleshooting DNS.

Summary
2:59-3:19
Ping -- right away if I can ping by IP address and not ping by name, I know it's DNS. NSLookup, if I really need to get into what the DNS server has or doesn't have in terms of records, and ipconfig to troubleshoot with the DNS cache, clear the cache, or get the client to check-in with DNS and re-register its records.

Using Ping and NSLookup
0:00-0:08
In this video, we're going to take a look at troubleshooting DNS, and I do all my troubleshooting of DNS from the Command prompt.

The Problem
0:09-0:27
Let's say I need to ping a particular computer. As soon as ping comes back with text like that, then I know it's a problem with DNS. Even if the computer didn't exist or wouldn't respond, I should still get an IP address back.

Connectivity or DNS Record
0:28-0:34
The next question in my mind is, is it a problem with network connectivity, or is it it a problem with that particular DNS record?

Ping the Active Directory Domain
0:35-0:59
What I just do is ping something else. That's the easiest way. The great thing about pinging the Active Directory domain is, not only does it tell you that DNS is responding, but if you want to get logged in, you need to be able to ping the name, so anytime you're having an issue where it comes back and it says, "Login server is not available". try to log in with the local account. See if you can ping the domain name. If you can't, that's what your problem is, because it can't locate a domain controller.

Who is My DNS Server
1:00-1:51
If we pretend that this ping northsim.com didn't return an answer, my next step would be to look at who is my DNS server? You can only see the DNS server using ipconfig /all. In my case, my DNS server is myself, 127.0.0.1. But, if that came back with an actual IP address, my next step would be to try to ping that IP address. If the DNS server responds to a ping on its IP address, but not a request for name resolution, then the problem is inside of DNS itself. I should restart the DNS service on the server; something like that.
We're trying to narrow down where the problem is. Between ping and ipconfig, I can do 99% of my DNS troubleshooting. In this case, I can only get so far. I've determined that, in fact, I can contact DNS. DNS is working, but I'm not getting an answer for that particular record.

nslookup
1:52-2:42
If you need to really dig into DNS troubleshooting, the command that you can use is nslookup. nslookup has two modes; one mode where you can just type the whole command, the other is an interactive mode. I'll show you each of those.
I would try next an nslookup, and my server tells me, nope, I can't find that client. It's a non-existent domain. That means my computer doesn't know about that domain. If I want to, I can go into the interactive mode, and you just type nslookup, hit Enter, and now I've got this different prompt. Anything that I type is prefaced by nslookup, and I can go in and run commands and say, hey, I'm only interested in email records. I'm only interested in this or that. I can be as discreet as I want to.

ls On the Domain
2:43-3:51
I'm actually going to do an ls eastsim.com. I want to actually have it list everything that's in the parent domain. Be aware, if you do an ls on the domain, what the computer's actually doing is a zone transfer. If you don't have rights to do a zone transfer, this particular command would fail. It's not going to be terribly useful, but because I'm poking around in my environment, I should be fine. It comes back and it says, well, I know about eastsim.com. The DNS server for that is dc1.northsim.com, but that's it. It doesn't know anything more than that.
Sales is a sub-domain, so either it's kept in the same file, which in this case, it isn't, or I'm probably missing a delegation. We'll go back into DNS and add the missing delegation. Now, if I do my ls eastsim.com, it's aware of the sales sub-domain, so I should be able to resolve that original record. In fact, I get an answer.

Non-Authoritative
3:52-4:15
You can see, it's telling me, it's a non-authoritative answer, which means that my DNS server, which is DC1, is not authoritative for sales.eatsim.com, but it got me an answer using the delegation. In fact, it's member2 that's authoritative for that particular domain. It doesn't usually make a big deal whether it's authoritative or non-authoritative. We just want an answer, but that's what it means, if you're wondering about that.

Summary

4:16-4:45

That's how we troubleshoot. Look for ping. Make sure that I can ping by name. If I can't ping by name but I can ping by IP address, that's definitely a DNS problem. Make sure I have connectivity with my DNS server. Try pinging my domain. If I really need to dig in and I can't get through with ping or ipconfig, then nslookup is a command that's just for working with DNS. It's a very extensive command. Again, try for the simple stuff first, the easy answers. If you can troubleshoot DNS with ping and ipconfig, hey, that's even better.

5.6.3 Clearing DNS Cache

Clearing DNS Cache
0:00-0:04
In this video, we're going to take a look at how to clear the DNS cache on both the client and the server.

Displaying the DNS Cache
0:05-0:46
Before we clear it, we should probably go and look at the DNS cache. We need to open up a command prompt, and I can display the DNS cache using ipconfig /displaydns. We can see that right now there's just one entry in there; workgroup1, which is actually mapping this name to 192.168.1.50.
If we look at /displaydns, I can now see that member2 has been added; as soon I pinged it; it had to go out and resolve that name to an IP address, and it was added to the cache.

Clear the Cache
0:47-1:19
If I'd like to clear the cache, it would be ipconfig /flushdns, and we can see that member2 has left the cache.
You might be thinking, "why did workgroup1 stay behind when I flushed the cache?" The only reason an entry would stay in the cache even though you flushed it is if that entry is in the hosts file. We'll take a look at the hosts file.

Hosts File
1:20-2:09
I'm going to run it as an Administrator. I'm in the C:\Window\System32\Drivers\etc., and if I display All Files, there's hosts, and you can see that there's an entry at the bottom. If I delete that entry, I should now be able to completely clear the DNS cache. I do a /displaydns and there's absolutely nothing in there.
That's how you work with the client cache, and in general, the only time you're going to be getting in there and doing that is if the server has changed its IP address, and we need the clients to be up to date right away.
Even when entries are stored in the cache, they have a certain time to live. When that expires, they would be taken out of the cache automatically.

/flushdns
2:10-2:50
You're only doing a /flushdns if you've changed something very recently and you know the entry is not expired from the cache as yet. It's something good you can try.
The other place where you'd want to try this is, if you try to contact a server or client that's not registered with DNS and DNS says No, there is no member2, that negative answer also gets stored in the cache.
If you then say, "oh, I've determined the problem is member2 is not registered with DNS", then you go ahead and register it. You would need to clear the cache in the client for that as well. You can register the IP address using ipconfig /registerdns.
If you need to clear the server cache, that's a different command.

Clear the Server Cache
2:51-3:09
Anything working with DNS is dnscmd, and the switch is /clearcache. If you think the issue is that you're getting the wrong answer from the DNS server, and it's the DNS server's cache, you'll want to run dnscmd /clearcache on the server.

Active Directory Records
3:10-3:26
If the problem is Active Directory records, /registerdns will not fix that from the domain controller. You need to restart the netlogon service.
That's how you manage the DNS cache, and that's how you clear it if you need to get an entry out of there right away. Otherwise you can just let it expire naturally and there should be no problem.

5.6.4 DNS Troubleshooting Facts

Use the following tools to troubleshoot DNS:

Tool	Description
Ping	You can use **Ping** in the following ways to determine if the problem is with DNS: • Ping a FQDN. If you get a message back that the hostname could not be found, the problem is DNS. If the problem were TCP/IP, you would get no response or a response that the destination host was unreachable. • You can determine if the problem is with your client contacting DNS or if it is a problem with DNS in the following ways: Ping the IP address of the DNS server for the domain you are in. If you can successfully ping the server, there is a DNS problem. Ping the IP address of the host you are trying to contact from another host. If you can ping by IP address, but not by the FQDN, there is a DNS problem.
NSLookup	**NSLookup** is a command for querying DNS. Using **NSLookup**, you can specify: • The DNS server to which the query is sent. • The types of records to view. • Other variables in the DNS record. Be aware that some NSLookup queries require permission.
IPConfig	IPConfig is a command you can use to display TCP/IP network configuration values and change DNS settings. Parameters you can use to troubleshoot DNS are: • **/displaydns** displays the content of the DNS cache • **/flushdns** clears the DNS cache • **/registerdns** causes the host to register its address with DNS

6.1 File Access

As you study this section, answer the following questions:

- How are NTFS permissions used to control access to file and folders?
- Why should you assign permissions to groups rather than user?
- How do logged on users get updated permissions?
- How would you determine the effective permissions when both NTFS permissions and share permissions apply?
- What is the difference between *explicit permissions* and *inherited permissions*?
- What will happen if you move a file that has explicit NTFS permissions to a different folder on the same NTFS partition or if you copy the file to a different folder on the same NTFS partition?
- What is the difference between a *soft quota* and a *hard quota*?

After finishing this section, you should be able to complete the following tasks:

- Configure NTFS permissions.
- Remove inherited permissions.
- Enable quota restrictions.
- Create a quota entry.
- Modify quota limits.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  - Manage NTFS Permissions
    - Configure NTFS Permissions
    - Configure Inherited Permissions
  - Manage Combined NTFS and Share Permissions
  - Configure Quotas
    - Manage Quota Restrictions
    - Create Quota Entries
    - Configure Quota Limits

This section covers the following 70-410 exam objective:

- 201 Configure file and share access.
  This objective may include but is not limited to:
  - Create and configure shares
  - Configure share permissions
  - Configure offline files
  - Configure NTFS permissions
  - Configure NTFS quotas

6.1.1 NTFS Permissions

NTFS Permissions
0:00-1:12
We're going to talk about NTFS permissions. The first thing I want to say about NTFS permissions might seem really obvious, but they're only set up on NTFS volumes. A particular drive has to be formatted with NTFS in order for me to be able to go into the Properties of a file or folder and see the Security tab. NTFS permissions are either Explicit or Inherited. In an ideal situation, when I set up my file servers, what I want to do is go ahead and create a system that's going to perpetuate itself. I would like to go through and create a Public folder that everybody has access to; some folders for the Departments, and then maybe individual folders so that each person has a public place, semi public place, and a private place to store data. Real life it's hardly ever this simple, but that's the goal. Any permissions that I set at Departments would be inherited at HR.
The reason that inheritance is great is, I can go in -- and really I would be setting up my permissions at HR -- and then any subfolders that are created by those users are going to get those permissions. When I set up my permissions, I'm going to grant the permissions to groups. Maybe I have an HR domain local group, and I'll give it modify. We're going to talk about the rights in a minute. Then they can go ahead and create whatever subfolders they need, and those rights will stay in effect.

Explicit or Inherited Permissions
1:13-2:04
NTFS permissions can be Explicit, which means I actually go out to the folder and assign those permissions, or they can be Inherited, which means I assign them at a parent folder. The subfolders inherit the permissions that were set on the parent folders. If I don't like that, inheritance can be disabled.
You want to be careful with that. If you're disabling inheritance every day, there's some type of a structural problem with your file system. If you need to for a particular folder, you can disable it and then reset the permissions. If I have an Inherited permission for that user or group, I've got to disable inheritance before I can set an Explicit permission. Inheritance can also be forced, so if I've messed up the security on subfolders, I can go up to the parent and say, "force inheritance right on down the folder structure and reset all of the NTFS permissions on the folders below this so I can get back to a point where I know where I'm at with my permissions".

Cumulative
2:05-2:49
NTFS permissions are Cumulative, so my effective permissions, meaning what I can actually do, is going to be the most permissive of all of the permissions that I have.
The only exception is Deny. If I set up a Deny permission, that's going to override any Allow permissions. Explicit permissions also override Inherited, but that hardly ever comes into play, because if I'm going to explicitly set a permission for that user, I'd have to break inheritance. If they've inherited something through a group and then I explicitly grant them a right, that Explicit permission technically will override the Inherited. That's the only spot where you would see an Allow permission override a Deny. For practical purposes, just keep in mind that Deny overrides Allow, and 99.9% of the time, that's exactly what's going to happen.

Effective Permissions
2:50-3:24
Let's say I am a member of this HR domain local group which has Modify, but I also happen to be a member of the IT group and that has Full Control. My effective permissions are the combination of these two. Since Full Control includes Modify, I'm going to get Full Control. I'm going to get the most permissive. If you're not sure what somebody's effective permissions are, you can go into the properties of the folder, and there's an Effective Access tab. They've changed the name of that tab with Windows server 2012. You can add them in and see what their permissions are actually going to be.

Basic NTFS Permissions
3:25-3:28
Let's take a look at the types of permissions that we have access to with NTFS.

Read
3:29-3:33
Read permissions allow the user to open the file or folder. They can read it. They can't make any changes.

List Folder Contents
3:34-3:39
List of Folder Contents applies only to folders and it lets them see the contents of the folder.

Read and Execute
3:40-3:49
Read and Execute lets them open the file or folder and execute programs in that folder. The difference there is that when you execute a program, it makes a temporary file, and that would be permitted if they have read and execute.
Write
3:50-3:52
Write "let me save content to the file or folder".
Modify
3:53-4:00
Modify is the same as Read, Write, and it also includes the ability to delete files or folders within that folder.
Full Control
4:01-4:10
Full Control is Modify, so everything above this on the list, Modify, plus the ability to do security functions, which means I can modify the NTFS permissions and take ownership if I need to.
Choosing a Permission
4:11-5:57
Let me say a few things about these permissions. First of all, in practice, we're really only given out three levels. Read and Execute means that they can get into the folder or the file and look at the contents but not change anything.
There's no point in giving out Read without Execute. Giving out Execute is not a security issue. You also would not give out Read without also including List Folder Contents; kind of mean. You can read this file if you can guess what it is. We wouldn't do the opposite either. We wouldn't List Folder Contents but not give them Read. So in practice, we just click Read and Execute. That means they can go in and look, but they can't touch. The next level up that I would grant is Modify. Modify means that they can save content and they can delete content. The only difference between Read and Execute plus Write versus Modify is that they have the ability to delete.
While there may be situations where you do want to give out read and write, hopefully it's not often, because I'll almost guarantee you, the minute you give out Read and Write and not Modify, about three in the morning that night you're going to get a call, come into work so that they can actually have you delete a file that they accidentally saved up there. Generally we just give out Modify. Full Control is usually reserved for network administrators because it allows me to handle the security of that file or folder. I can go into the Security tab and adjust the NTFS permissions. I also can take ownership. The owner has the inalienable right -- which means you can't take that right away from them -- to change the security of that file or folder. If somebody got themselves locked out of a file or folder, as an administrator I could come in, take ownership. Now that I'm the owner, I have the ability to change the permissions, and I can reset it so that we can get back in. We give out Read and Execute if they can read but not touch, Modify if they can do all the normal functions, saving Full Control back for the administrators. These are my basic NTFS permissions.
Advanced NTFS Permissions
5:58-6:29
There actually are Advanced NTFS Permissions, but these are seldom used. We try to work with the basic permissions as much as we can. I've been in the industry about 18 years, and I can't think of one example where I had to go to advanced NTFS permissions and I couldn't just use the basic permissions.
New with Windows server 2012, Advanced permissions allow the administrator to set a condition for the permission. I'm sure that there are some advantages to this, but generally speaking, if you set up the basic permissions correctly, you shouldn't need to set a condition for them to get into that particular folder.
Best Practices
6:30-7:02
Let's talk about best practice for NTFS permissions. Assign permissions as high up in the folders structure as possible. We know inheritance is going to make them flow down. I don't want to be giving in permissions, taking them away, giving them in; I just want to sign the permissions at one folder. Whatever the users do, they'll just flow right on down using inheritance. Assign permissions to groups, not individual users.
If I assign the permission to a group, then if that person leaves the company, I just pull that account out of the group and I add the new person into the group. If I assign it to a user, I'm going to have to revisit the share.

Special Identities

7:03-7:54

I can also use special identities. Special identities are groups that are created by Windows. Administrators cannot control the membership of these groups. You get into these groups by virtue of what you do. For example, one of the special identities is the everyone group. Everyone is in the everyone group, just as a matter of default. There's another special identity, "authenticated users". It's anybody who's logged into the network. We use anonymous users for websites. That's anybody who hasn't logged into the network.

There is a network, "special identity", which is anybody who's accessing that file across the network, versus an interactive special identity, which is anyone who has sat down at the keyboard and logged in. We use these special identities when we want to grant permissions to people based on what they're doing. I don't know in advance exactly who that's going to be, but I know anybody who sits down at the keyboard and logs in should have this NTFS permission.

Summary

7:55-8:17

NTFS permissions are great for security. It lets us lock down files and folders. I want to assign my permissions as high up in the folder tree as possible, counting on inheritance to have those permissions flow down. I can break inheritance if I need to, but I'm going to avoid that if I don't need to, and they affect everybody all the time. When you're talking about securing files or folders, NTFS permissions are really the way to go.

6.1.2 NTFS Quotas

NTFS Quotas
0:00-0:12
Let's talk about NTFS quotas. First of all, it may seem obvious, but NTFS quotas can only be set on NTFS volumes. You'd have to have that drive formatted with NTFS in order to be able to change the quota.

Quota
0:13-0:27
Quotas limit the amount of space that users can consume on that volume. The amount of space that they're using is tracked by ownership. Whatever files they create, they're the owner of those files, and those count against their quota.

Different Limits
0:28-0:43
If I have particular users that should have different limits than everybody else, I can make a quota entry and give them a different limit. The boss comes in and says, "Everybody else gets 500 MB, but if you don't give me a TB, you're fired." I can certainly do that.

Hard Quota
0:44-0:54
When I set up my quotas, I can set either a hard quota or soft quota.
A hard quota will cut the user off once they've consumed their allotted amount of space.

Soft Quota
0:55-1:18
A soft quota just sends the administrator a message by putting it in the event log when that person has met their quota but it doesn't cut them off. Users will expand their data to fit the space available. You can never have too much hard drive space. If you set a soft quota, you're really not going to be doing anything effective. You need to limit them if you want to actually control the amount of space that they're using.

Warning Limit
1:19-1:33
You can also set a warning limit so that when they hit a certain percentage of whatever they're allotted, they would get a warning.
NTFS quotas affect all the files that that user saves to that particular drive.

File Server Resource Manager
1:34-1:47
If you're looking for more flexibility with quotas, you want to have them be on a per folder basis, something like that. You've really got to get into File Server Resource Manager, which will give you a lot more options. NTFS quotas are great if you just need to limit the users on a per volume basis.

6.1.3 Configuring NTFS Permissions and Quotas

Configuring NTFS Permissions and Quotas
0:00-0:04
In this video, we're going to take a look at configuring NTFS permissions and quotas.

NTFS Permissions
0:05-0:19
NTFS permissions affect all the users all the time, so even though you have to share out folders for users to be able to get to data over the network, you want to open up your share permissions wide open and then lock it down on the NTFS side.

Creating a Group
0:20-1:17
Let's go through and just take a quick look at some groups that we can use. In Sales, I have four users, Beth, George, Mike, and Sally, and I have a global security group called Sales Users, and they're all members of that group. We want to use global groups to organize our users. We like to use domain local groups to organize permissions. What I'm going to do here is, make a domain local group and add my global group into it, just to make sure I'm following Microsoft best practice. It needs to be a security group in order to be able to add it to the Security tab. Now that I've created my group, I'm going to add in that Global group, so my users are a member of Sales Users. Sales Users is going to be a member of Sales Data, and Sales Data is the group that I'm going to use when we set up our NTFS permissions.

Adjusting Permissions
1:18-1:38
Now that we've got our groups set up, let's go in and take a look at how we would do this. I've created a folder on drive (C:) for departments, and then I'm going to go in and create a folder for the Sales Department. To adjust the permissions, we're going to right click and go to Properties, and then go over to my Security tab.

Inherited Permissions
1:39-2:21
NTFS permissions are inherited from whatever the parent folder is. So, the permissions I'm looking at right now are the default permissions for the C: drive on my server, and that's being inherited from the C: drive to Departments, from Departments to Sales. If I just want to add something in, I would hit Edit and add in whatever group I want. The problem is going to come in because by default on the root of C, users have the right to read that. You see how these check marks are kind of grayed out. That means they're inherited permissions. I can add in new groups, but I can't take away those inherited permissions unless I break inheritance. So we're going to go into Advanced, and we'll disable inheritance.

Disabling Inheritance
2:22-3:00
Now it says you've got a bunch of permissions in there. You want to copy them, convert them into explicit permissions, or just want to remove everything? I recommend you convert and then take away what you don't want. That way you don't accidentally end up a folder with no rights on it at all. Even if you did, you could fix that.
Now I go to Edit -- these are Explicit permissions -- and I can remove them. We'll go ahead and add that sales data, and we'll give it the Modify permission.

Types of Permissions
3:01-4:14
These permissions here are basic NTFS permissions. You can see we've got Read, List folder contents, Read and execute. Read and execute really is the bottom line permission I would give out. I wouldn't bother giving these two separately, and if we uncheck everything, you'll notice if I click Read and execute, it pulls in List folder contents and Read with it as well. Read and execute plus Write, the only difference between that and Modify is, Modify includes the ability to delete. Generally that's what we give out. If you're giving Read and Write but not Modify, almost guaranteed you're going to quickly be getting a call from somebody saying, well, I accidentally saved something into that folder. Can you delete it for me?
If you're concerned about people deleting other people's stuff, then you could go in and grant Modify to CREATOR OWNER, meaning if I'm the CREATOR OWNER I've got Modify, which includes delete, and then the Sales Data people just have Read and Write for whatever they want, but generally we just give out Modify to the group and we're happy.
So I've got my permissions set up.

Creating a Subfolder
4:15-5:16
I'm going to hit OK. If I create a subfolder, the subfolder's going to inherit those permissions. Let's do that and take a look. You can see it's still got SalesData. They've got Read and Write, CREATOR OWNER. You've got Special. The effective permission should be Modify. If I don't like that, I can change it. Now, what I've just done is added permissions further on down the line. That's not the best way to handle NTFS permissions. On a good day, I want to assign permissions as high up in the folder tree as I can and walk away. If I've got permissions that get added in at one folder and taken away in another folder, not only am I kind of teasing people saying, well, if you can hack into this you could have it, but on top of that, it's going to make it difficult for somebody else to troubleshoot.
If I do have a question as to what somebody's effective permissions would be, we can check that as well, using the operating system.

Effective Permissions
5:17-6:16
We go back into Properties, Security, and when I go into Advanced, you'll notice that we have an Effective Access tab where I can go in and specify a particular user. Right now this is going to evaluate George based on his current group membership. If I wanted to see what would happen if I add them to a group, I could specify an additional group that they're not a member of quite yet and take a look at that as well, but let's just see what he's got. So based on my NTFS permissions, he's picking up almost everything. He can't Delete subfolders and files, but he's got overall Delete. He just doesn't have it on the File Permissions. He doesn't have any of the security stuff because we really just gave out Modify.
If I've gone through and I have added and taken away from the NTFS permissions further down the folder tree and then I get to the point where I say, something's messed up, I really just want to roll it back to what it would look like if I'm just doing inheritance.

Rolling Back Permissions to Inherited
6:17-7:35
You can do that. We'll go back up to Departments. We've made some changes at data. We're not happy with those changes. We just want to go back to all the subfolders have whatever's going on at sales. Right click Sales. Go to Properties, Security, Advanced, and this check box forces Inheritance down the folder tree, assuming that I have the right to do that. So, it's going to Replace all child object permission entries with inheritable permissions from this object. It's going to get rid of all the explicitly defined permissions. Am I okay with that? Yes, absolutely. That's exactly what I'm trying to do, and that has gone through, changed the permissions at Data. So if I go into Data we should not see the users inside there, and in fact we don't. So we forced inheritance right on down the line.
There's not much else to NTFS permissions. The only other thing that you can really adjust is, if you go into Advanced, these permissions here are NTFS advanced permissions.

Advanced Permissions
7:36-8:14
I can go in and hit either Add or View, and then I could set up different permissions if I click Show advanced permissions. Generally, we don't work with the Advanced Permissions. They've been there since I've been working in the field, and in 18 years I've never had to do this, but it's possible that you may have to come in and do the Advanced Permissions, so I want to make you aware of them.

Conditions
8:15-9:15
I also can set up conditions as well. I could add a condition and say, well, if the User is a member of a particular Group, a Member of any, a Member of each, and then I could add items in here. Here, whether they're a member of either SalesData or SalesUsers, whatever I set up for permissions will go for that particular condition. Generally speaking, you should be able to set up the NTFS permissions without using conditions, but it's a new feature of Server 2012, and it might be a little bit more effective if I could go in and say, well, as long as the user is not a member of any of these groups, in that case, they can have these permissions. Instead of having to go in and deny those groups, I could actually set up a condition instead of doing a deny permission.

Cumulative Nature of Permissions
9:16-9:42
Permissions are cumulative, so users inherit permissions from whatever group they're a member of, and we just saw that with George. He's a member of SalesUsers, which is a member of SalesData, and that's where his permissions are coming from. The only thing that would override that would be a Deny permission. I could be a member of 40 groups that have Full control. If I'm a member of one group that has Deny, my effective permissions will be Deny.
On the topic of NTFS, we'll take a look at NTFS quotas before we finish up.

NTFS Quotas

9:43-11:15

NTFS quotas are set on the entire volume. Let's take a look at volume (E:). I want to right click it and go to Properties. Then I'm going to click on Quota. I would need to Enable quota management and then Limit the disk space to whatever I'm going to limit it to. Probably 1 KB is a little bit too small. You can set up a warning for the users. At this point, what it's going to do is, it's going to take a look. It's going to give them a warning when they get up to 275 MG, and then I would come in here and say, "send a message to the event log when they hit their warning level or their quota limit". It's not going to cut them off unless I come up here and check "Deny disk space to users exceeding quota limit". If I don't check that, it's not going to do much of anything.

With NTFS quotas, the quotas are calculated based on ownership. So even if they had four shares they could get into on this E: drive, all the files that they save in any of those four shares are all going to count against their quota, so you probably only want to use NTFS quotas when you're making a volume; let's say, that has individual user directories. If you're looking at department shares, you would want to use File Server Resource Manager, and you have a lot more flexibility with quotas using that utility.

Quota Entry

11:16-11:39

If I do have a particular user that I want to give more space to, I can also make a Quota Entry. By default, administrators have unlimited, but I could come in here, make a quota entry, and say George has absolutely no limits, or I could put a specific limit just for that user.

Summary

11:40-11:53

So, NTFS quotas restrict users to how much space they can use. It's done on a per-volume basis. It's tracked by ownership, and if I need to give somebody a different quota I can do a quota entry. That's how you manage NTFS.

6.1.4 File Access Facts

Permissions are assigned to resources and not to users or groups. The two types of permissions are:

- NTFS permissions control access to folders and files stored on an NTFS partition.
  With NTFS permissions, each file and folder has an access control list (ACL).
  The ACL identifies the users or groups and their level of access to the folder or file.
  NTFS file permissions are available only on NTFS volumes or partitions.
  NTFS permissions are in effect when files are accessed through the network or when they are accessed locally.
  The two types of NTFS permissions are:
  - Standard permissions
  - Special permissions
- Shared folder permissions are assigned to a shared folder. Key facts about shared folder permissions are:
  Shared folder permissions are in effect only when the resource is accessed from the network. For example, denying access using Shared folder permissions will have no effect on the user's ability to access files when the user logs on locally. In that case, only the NTFS permissions will control access.
  When both share and NTFS permissions apply:
  - You determine the effective permissions of each type using the most *permissive* permission.
  - You then compare the effective permissions of both NTFS and share permission.
  - The more *restrictive* of the two sets of permissions takes effect.

The following table summarizes the permissions for folders and files.

Permission	Allowed Actions
Read	View folder details and attributes. View file attributes; open a file.
Write	Change folder or file data and attributes.
List Folder Contents	Includes all Read actions and adds the ability to view a folder's contents.
Read & Execute	Includes all Read actions and adds the ability to run programs.
Modify	Includes all Read & Execute and Write actions and adds the ability to add or delete files.
Full Control	Includes all other actions and adds the ability to take ownership of and change permissions on the folder.

When setting up or managing NTFS permissions, be aware of the following concepts:

Concept	Description
Ownership	Ownership affects access and assigning permissions as follows: • Every object, including files and folders, has an owner. • The owner is typically the user who created the file. • The owner has full control over the file and can assign permissions to the file. • Administrators have the **Take Ownership** right to all objects. Administrators can assign ownership of a file or folder even if they do not have permissions to access the file. • You can reassign ownership of a file or folder to give a user all permissions. You might reassign ownership when someone leaves your organization. • If you cannot access a file because of insufficient permissions, take ownership of the file and modify the permissions.
Explicit vs. inherited permissions	Permissions are also called Access Control Entries (ACE). An ACE can either allow or deny access, and can be configured explicitly or inherited. • *Explicit* permissions are set on the object; *inherited* permissions are set on the parent object and apply to the contents of the folder. By default, when new files or folders are created, they inherit the permissions of their parent folder. You can block inheritance by deselecting **Allow inheritance** in the NTFS permissions window. When blocking inheritance, a recommended practice is to copy the inherited permissions, so you will have a record of the inheritable permissions. If you need to reset the inherited permissions for a file or folder, select the parent folder and then select the **Replace the permissions of all existing child objects** option under the **Advanced options** of the **Security** tab. Removing inheritance is an advanced NTFS permission option. • The *allow* permission grants the user, group, or computer the specified permission to the object. • The *deny* permission restricts access to the object. The deny permission overrides the allow permission, unless the deny permission is inherited and the allow permission is explicit. Explicit permissions take precedence over inherited permissions, even inherited deny permissions. Use the deny permission only when you want to override specific permissions that are already assigned.

	• Permissions are cumulative. Users gain the sum of all permissions granted to the user account and any groups. In Windows Server 2012, you can check the effective permissions for a file or folder on the **Effective Access** tab. The permissions shown in the **Effective Permissions** tab are approximate permissions, and can vary depending on how a user logs in or how they access the resource.
Copying or moving files	You must have the following permissions to copy or move a file: • To copy a file or folder, you must have Read permissions to the source file and Write permission to the destination location. • To move a file or folder, you must have Read and Modify permission to the source file, and Write permission to the destination location. Copying or moving files or folders that have NTFS permissions assigned can affect the permissions on the file or folder. • If you copy or move a file to a non-NTFS partition, all permissions are removed. • If you copy or move a file to a different NTFS partition, the file will inherit the permissions assigned to the parent partition and folders. • When a file has explicit NTFS permissions assigned to that file: If you copy or move the file to a different NTFS partition, the explicit permissions will be removed. If you move the file to a different folder on the same NTFS partition, the explicit permissions will be kept. If you copy the file to a different folder on the same NTFS partition, the explicit permissions will be removed. In all cases, the file will also inherit permissions from its new partition and folder. • Use the **robocopy** and **xcopy** command line utilities to copy files while maintaining the NTFS permissions (even when copying between partitions).
Using **icacls**	Use the **icacls** command to manage standard NTFS permissions from a command prompt. Be aware of the following switches: • **/grant** grants the specified user access rights. • **/deny** explicitly denies the specified user access rights. • **/save** saves and enables the ability to restore the user access rights. • **/restore** restores user access rights.

Special permissions allow granular (i.e. very specific) configuration beyond the six standard NTFS permissions. The following table illustrates how the special permissions correlate with the standard NTFS permissions:

Special Permission	NTFS Permission					
	Full Control	Modify	Read & Execute	List Folder Contents	Read	Write
Traverse Folder/ Execute file	X	X	X	X		
List Folder/Read Data	X	X	X	X	X	
Read Attributes	X	X	X	X	X	
Read Extended Attributes	X	X	X	X	X	
Create Files/Write Data	X	X				X
Create Folders/Append Data	X	X				X
Write Attributes	X	X				X
Write Extended Attributes	X	X				X
Delete Subfolders and Files	X					
Delete	X	X				

Read Permissions	X	X	X	X	X	X
Change Permissions	X					
Take Ownership	X					

Be aware of the following special permission details:

- Use special permissions to determine the level of permissions propagation, such as applying to all files and folders and subfolders, or to only the files in the folder.
- Special permissions offer finer control over the actions that can be performed on the file or the folder. To edit these permissions, click the **Advanced** button on the **Security** tab in the file or folder properties.
- Permissions are cumulative. If you are a member of two groups, both with different NTFS or special permissions, you will have the combined permissions of both groups (known as *effective permissions*).
- In Windows Server 2012, you can set a condition for a special permission.

Best practice for permissions include:

- Assign permissions as high up in the folder structure as possible.
- Assign permissions to groups, not individual users. You can use special identities which is a group created by Windows.
- Use domain groups to set permissions.
  Set the Group scope as Domain local
  Set the Group type as Security

NTFS quotas limit the amount of space that a user can use on an NTFS volume. Be aware of the following regarding quotas:

- Quotas are tracked based on file ownership.
- A quota amount applies to all users in the group.
- Quota entries can be used to specify a different limit for a designated user.
- If you use a *soft quota*, the administrator is notified when a user meets the quota limit.
- If you use a *hard quota*, a user is not allowed to use more disk space.
- You can set a *warning limit* that notifies the user when a specified percentage of their quota limit is reached.
- File Server Resource Manager provides an administrator more flexibility by allowing quotas on a folder basis.

6.2 Access-based Enumeration (ABE) and Volume Shadow Copy (VSS)

As you study this section, answer the following questions:

- How does using access-based enumeration on shared folders modify what users can see?
- What tool is used in Windows Server 2012 to enable ABE?
- How should client work patterns affect your shadow copy schedule?
- How much disk space do shadow copies take by default?
- What is the maximum number of shadow copies the system can store and what happens when the system reaches this limit?
- What happens to NTFS permissions when you restore a file or when you copy a file?
- Why is it recommended that you place shadow copies on different volumes?

After finishing this section, you should be able to complete the following tasks:

- Configure ABE.
- Enable shadow copies on a volume.
- Configure shadow copy settings, including storage location, size, and schedule.
- Create snapshots.
- Save, copy, or restore previous versions of files.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  Configure Volume Shadow Copy Service (VSS)
  - Enable and Configure Shadow Copies
  - Restore Previous Versions of Files and Folders

This section covers the following 70-410 exam objectives:

- 201 Configure file and share access.
  This objective may include but is not limited to:
  Configure access-based enumeration (ABE)
  Configure Volume Shadow Copy Service (VSS)
- 302 Create and configure virtual machine storage.
  This objective may include but is not limited to:
  Manage checkpoints

Access-based Enumeration (ABE) and Volume Shadow Copy (VSS)
0:00-0:07
We're going to talk about access- based enumeration, which they abbreviate ABE, and Volume Shadow Copy, which we abbreviate VSS.

Access-based Enumeration (ABE)
0:08-0:41
Access-based enumeration really is exactly what its name says it is. To enumerate something is to look at the contents. If it's access-based enumeration, what it means is that the user can only see the files and folders to which they've been granted access. If they go into a folder that has ABE turned on and there are sub folders or files to which they don't have access, then they won't even see them. This is a great feature, because if you let somebody in a folder and then you lock them out of the sub folders, I can almost guarantee you they will spend the rest of their natural born life trying to get into that particular folder.

Volume Shadow Copy (VSS) History
0:42-0:48
Volume Shadow Copy is a technology that has been improving with every generation of Windows Server.

The Problem Solved by VSS
0:49-2:15
But essentially, the initial problem that we had was this: backup software could not back up open files, because those files were locked. When I worked as a young LAN administrator using NT40, we'd get tickets to the help desk saying "my most important file in the world, it's so critical, I've never closed it in six months, I've got to have that restored from backup," and we'd have to very gently explain to the user, "I'm so sorry but we haven't had a good backup for six months because you've never closed it." The initial problem that we wanted to resolve was the ability to back up files that were open. What Volume Shadow Copy does is take a snapshot of that file, whether it's open or not, if it's changed, and it saves the difference in a special area.
Originally what that allowed us to do was, if the file was open, we could back up the snapshot and as an aside. Anytime you hear the word snapshot, you want to be thinking some flavor of Volume Shadow Copy. This progressed into saying "if we're taking a snapshot of the file, why not use it for other things besides the backup service?" It became integrated into the file system to allow users to go back in and restore a file that's changed from a previous snapshot. We'll do that on the previous version's tab. Then from there, it morphed into being able to take snapshots of Active Directory, and go in and restore pieces of Active Directory from the snapshot as well. There are lots of places in the operating system where we would see Volume Shadow Copy in action. We're really going to focus on how it works in the file system.

Volume Shadow Copy
2:16-2:37
Volume Shadow Copy takes snapshots of files whether they're open or not, and these are files that have changed. That's going to come up a little bit later on in this presentation again. We turn it on, on the Shadow Copies tab in the Properties of the drive. They're enabled on a per volume basis. What that means is, if you have different volumes with different Volume Shadow Copy requirements, you've got to partition the drive.

Settings of VSS
2:38-2:40
In the Volume Shadow Copy Settings, we have some options.

Limit Space
2:41-2:46
The first thing we can do is limit the space used for shadow copies. Maybe we don't want these things to take over the hard drive.

Schedule
2:47-3:13
We can schedule when Volume Shadow Copies are made. By default, the operating system is set up to take two Volume Shadow Copies per day. Microsoft recommends that you not set it to be more frequent than every hour. It's only going to save as much as you have limited the space, so you want to play off how long you're going to have these shadow copies. If I make two a day, I probably can get about a month's worth of shadow copies in. If I'm making them every hour, I might only be able to roll back to a week ago.

Location

3:14-3:55

We can also change the location where the shadow copies will be kept. There are a couple reasons for this.

First of all, they can be moved for extra fault tolerance or to save space. Maybe I say, "let me move the shadow copies for volume C: onto drive D:. That way if C: goes down, I can restore from backup." That gets me back to maybe 8:00 a.m. in the morning, but if there was a shadow copy made at 1:00 p.m., that shadow copy will be on another drive, and I can even get drive C: all the way back to 1:00 p.m.--what all those files look like at that particular time. If you're going to move the Volume Shadow Copies, they must be moved before being enabled. Once you turn them on, you go into settings, and the spot where you move them is going to be grayed out.

Restoring VSS

3:56-4:21

To use Volume Shadow Copy, we restore our shadow copies using the Previous Versions tab.

This is important because people don't always use the right vocabulary. Some people say Volume Shadow Copies. Some people say shadow copies. Some people just say previous versions, because that's the tab where I go to actually restore the shadow copies. Any one of those terms would be correct. Remember, our Volume Shadow Copies take snapshots of files that have changed.

Previous Versions Tab Shows No Copies

4:22-4:49

If you go into the previous version tab and there are no copies in there, it could be one of two problems. First thing is, Volume Shadow Copy might not be enabled. By default, it's not enabled, so you won't have Volume Shadow Copies on a server unless you turn them on. With client operating systems, they're on by default for the C: drive, but not other drives.

In Windows Server 2012, nothing is on until you turn it on, or it might be that the file or folder has not changed. That would be another reason that that previous version tab would not show any copies.

Summary

4:50-5:04

Access-based enumeration lets me make sure that users can only see the files and folders to which they have access. Volume Shadow Copy is used to take snapshots of files that have changed whether they're open or not, which lets me go back in and restore those snapshots to roll the file back to a particular time.

Configuring ABE and VSS
0:00-0:06
In this video, we're going to look at Access-based enumeration (ABE), and volume shadow copy (VSS).

Access-Based Enumeration: ABE
0:07-0:28
Access-based enumeration, ABE, is exactly what it sounds like. They're going to be able to enumerate or see only those files or folders to which they have access. If they open up a folder and there's a subfolder in there, and they don't have rights, they're not even going to be able to see it. I would set that up using File and Storage Services. It's done on Shares, and we have some Shares here; set it up on Simple.

Set Up
0:29-1:15
I'm going to right click and go to Properties. Here are my Permissions. I can set up the Permissions in here. We're interested in Settings. Enable access-based enumeration, you can see right there, they will only be able to see the files and folders that they have permission to. If they don't have permission, Windows is going to hide it.
We also can come in here and Allow caching of the share or not allow it; very rudimentary; set up for offline files. We can even go through and set up that any access to the share is going to be encrypted.
The next thing we're going to look at is volume shadow copy.

Volume Shadow Copy
1:16-1:46
Basically what volume shadow copy does is take a snapshot of the file. This is done on a per volume basis. I've got to go into the Properties of the drive itself. We're going to set it up on the Shadow Copies tab. You can see, by default, it's Disabled.
If you want to make any changes to this, it's better to make the changes before you enable volume shadow copy.

Location of Shadow Copies
1:47-2:02
I'm going to click on Settings, and notice, I can change where the shadow copies are going to be stored. If I'm concerned about volume C: and I say well, if C: were to die, I want to make sure that those shadow copies are still available; then, I could store them on a different volume.
The other settings, like Maximum use limit, which is the maximum amount of space that can be used for volume shadow copies and the Schedule -- both of those can be adjusted after the fact.

Maximum Use Limit
2:03-2:17
Where they are located can only be adjusted before I turn them on.

Schedule
2:18-3:41
If we go into Schedule, we can see that default volume shadow copy is set up to take two shadow copies; one at 7 a.m. on every weekday, and one at noon on every weekday. Let's say that this is on a Tuesday. My backup finished at 8 a.m. Volume shadow copy is making a snapshot of everything's that changed at noon. If the server dies at 3 p.m. and if my volume shadow copies are located on the same drive, the best I can do is restore from backup and get myself back to 8 a.m.
If I store the shadow copies on a different drive, I could actually get the users back to noon by first restoring that drive from backup and then restoring the shadow copies.
Volume shadow copy can keep about 64 copies, so with 2 a day I've got about a month of data as long as I don't run out of space. You can set up whatever schedule works for you. You can see Microsoft does not recommend that you create these more frequently than once per hour. That seems very frequent to me.
I'm going to enable this with just the default settings; it says I'm going to use the default, sure. Notice it says too, if your server has a lot of activity on the disc, maybe you might want to make less shadow copies. This is going to put some stress on the disc, because whenever it makes the shadow copies, it's going to be doing a lot in the background; so keep that in mind when you set your schedule, how often and when it's going to do this. I would probably pick noon, thinking that everybody's at lunch; it's a good time to create them.
To go in and restore from a shadow copy, I just go to the parent folder of the file.

Restoring From a Shadow Copy

3:42-4:19

If I go into Properties, I'll have a Previous Versions tab. If this Previous Versions tab says No volume shadow copies, there would be two reasons for that; either I don't have it turned on, or nothing in the folder has changed, because volume shadow copies only takes a snapshot if the files are open or if they've changed.

If I wanted to bring the entire folder back, the Restore button does that, and it's actually grayed out for Windows, because it's a huge folder.

If I just want one file in that folder, it's better to copy it to another location and just get the one file that I want.

Summary

4:20-4:45

That's how we work with Access-based enumerations and volume shadow copy. It is abbreviated VSS, so anytime you see VSS or volume shadow copy, or shadow copies, that means we're taking snapshots of files that are open or that have changed, so that we can roll them back or we can back them up.

Access-based enumeration means I can only see those files or folders to which I actually have rights, and it's done only on the share. That's how we work with those two technologies.

6.2.3 ABE and VSS Facts

Access-based Enumeration (ABE) restricts users from seeing files and folders to which they do not have access when browsing content on the file server. ABE eliminates user confusion caused when users connect to a file server and encounter a large number of files and folders that they cannot access. ABE applies to domain-joined computers; it is not active when viewing files and folders in the local file system.

The differences in implementing ABE in WS2008 and Windows Server 2012 are shown in the following table:

Version	Description
Windows Server 2008	On a computer that is running Windows Server 2008: • ABE is enabled by default on every folder that is shared using the **File Sharing** feature. However, ABE is not enabled by default on the following types of shared folders: Shared folders that are created with Share and Storage Management, Advanced Sharing in Windows Explorer, or the **net share** command. Volumes. Folders or volumes that are shared for administrative purposes, such as C$ and ADMIN$. • ABE can be manually enabled or disabled on individual shared folders and volumes by using Share and Storage Management. This snap-in is available after a folder or volume has been shared. Share and Storage Management are accessed in the File Services server role in Server Manager and in Administrative Tools. You can also install it manually in Server Manager by adding the **File Server** role service to File Services.
Windows Server 2012	In Windows Server 2012, Server Manager replaces Share and Storage Management. To enable ABE: • Open File And Storage Services in Server Manager. • Open the share folder and select the folder for ABE. • Check the **Enable access-based enumeration** option in **Properties** > **Settings**.

Volume Shadow Copy Service (VSS) is a feature that automatically makes copies of user files at regular intervals. Enabling VSS allows you to:

- Recover deleted files or folders.
- Recover a previous version of a modified file.
- Compare a file with a previous version of that file.

Keep in mind the following about VSS:

- Shadow copies are enabled on a volume, not specific folders or files.
- Always perform regular file server backups. Shadow copy does not replace regular backups.
- Use the **Shadow Copies** tab to enable VSS and configure storage locations and schedules.
- Use the **Previous Versions** tab of a volume, folder, or file to view and manage previous versions.

You should know the following about implementing shadow copies:

VSS Area	Considerations
Scheduling	By default, the system takes two snapshots (shadow copies) of volume data each day (Monday through Friday). • You can modify the schedule to customize when and how often snapshots are taken. • You can also manually take a snapshot. • Base your VSS scheduling on client work patterns. If possible, schedule copies to occur during off hours. Schedule copies to occur more or less frequently depending on how often the data changes. • Do not schedule copies to occur more frequently than once an hour.
Storing	By default, up to 10% of the volume will be used for storing shadow copies. • The amount of disk space required for each shadow copy is typically less than the size of the current file. This is because shadow copy saves only incremental changes that have been made to each file, not the entire file (unless necessary). • Disk space usage for past copies can be customized by using either a percentage or a fixed amount. • At least 300 MB of free space must be available. • The system can store up to 64 shadow copies. • When no more disk space is available, or when the 64 copy limit is reached, the oldest copy will be deleted when a new copy is scheduled to be made. Once deleted, a shadow copy cannot be retrieved. • By default, shadow copies are saved on the same volume. It is best practice to place shadow copies on a different volume. Doing so improves performance and ensures that certain conditions will not affect the ability to save copies. Configure the copy location when you enable shadow copies to prevent losing existing copies. When configuring VSS storage, the following should be considered: • Do not enable VSS on volumes that use mount points or on dual-boot computers. You should not enable VSS on volumes that use mount points because the mounted drive is not included when shadow copies occur.

	• To allow defragmentation without causing previous versions of files to be deleted, format source volumes on which you plan to enable VSS with allocation unit sizes of 16 KB (kilobytes) or larger. If you plan to use NTFS compression on the source volume, do not use an allocation unit size larger than 4 KB, or you may lose older shadow copies faster than anticipated on very fragmented drives. • Before deleting a volume, disable VSS. If the volume is deleted first, VSS continues to run and will generate **Event ID: 7001** errors when the shadow copy event fails.
Recovering	You can recover a file, folder, or a volume. • Keep in mind the following when using VSS: Restoring files overwrites existing files. Restoring folders restores deleted files and overwrites existing files but does not delete any new files that have been added since the shadow copy was made. Restoring large directories has a negative impact on performance. If possible, restore individual files instead of directories. • You cannot revert a volume that contains system files. • When recovering a file, folder, or volume: Open a volume, folder, or file to view the contents of the previous version. The previous version is opened in read-only mode, so you cannot make changes. To make changes to a previous version, save a copy to a new location or with a new name, and then make the changes.
NTFS Permissions	NTFS permissions on previous versions depend on the action taken: • Restoring a file retains the file's permissions. • Copying a file to a different location sets the file permissions to the default permissions of the new location.
VSSAdmin	Use VSSAdmin to manage the Volume Shadow Copy Service from the command line. Be aware of the following options: • **list shadows** lists existing volume shadow copies. • **list writers** lists subscribed volume shadow copy writers. • **revert shadow** reverts a volume to a shadow copy. • **query reverts** queries the status of in-progress revert operations.

6.3 Shares

As you study this section, answer the following questions:

- What tools are available to create and manage shares?
- What is the effect of appending $ to a share name?
- What permissions do you need to share a folder or configure share permissions?
- What is the difference between the read permission and the change permission?
- What is the recommended strategy for assigning permissions to a share to provide the required effective permissions?
- What are the differences and similarities between NTFS permissions and share permissions?

After finishing this section, you should be able to complete the following tasks:

- Configure and manage shared files and folders.
- Restrict share access through share permissions and user limits.
- Enable and disable share caching.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  - Share Folders and Configure Share Permissions
    - Share Folders
    - Manage Shared Folders
    - Manage Share Caching
    - Configure Share Permissions

This section covers the following 70-410 exam objective:

- 201 Configure file and share access.
  This objective may include but is not limited to:
  - Create and configure shares
  - Configure share permissions
  - Configure NTFS permissions
  - Create and configure Work Folders

File Shares
0:00-0:25
Let's talk about file shares. In order for users to access data across the network, the folder needs to be shared out. There's a couple different ways of sharing it. We're going to take a look at Basic File Sharing, which is intended to really help you get things set up quickly and properly, and then we'll dig into Advanced File Sharing, which gives you more control over how the sharing is going to be done.

Simple File Sharing
0:26-0:43
Simple file sharing came in with Vista and Windows Server 2008, and then they modified it a little bit with 2008 R2 and Windows 7, but the great thing about it is it handles NTFS permissions and Share permissions, so you don't have to actually deal with them separately.

NTFS and Share Permissions
0:44-1:01
NTFS permissions are always in effect. Share permissions only affect people that access the data over the network. When they combine, you have to know what's going to happen. Simple file sharing takes care of that for you, so you don't have to worry about what the ultimate result will be.

Everyone Special Identity
1:02-1:15
It's going to assign the Everyone special identity the Allow-Full Control share permission, and that's fine; you don't need to be alarmed about that, because what you'll do is lock it down on the NTFS side.

Folder Name is Share Name
1:16-1:28
It's going to use the name of the folder as the share name, so you want to make sure that name is short, it doesn't have any spaces, because in the simple file sharing, you don't have control over what the share name is going to be.

Universal Naming Convention
1:29-1:57
When we share out folders, we need to know how to get the users into those shared folders, and what we use for that is the Universal Naming Convention. The Universal Naming Convention basically goes like this: \\SERVER\SHARE. So if this is FS1 and my share is data, it would be \\FS1\data, and I'm good to go. The nice thing about simple file sharing is that, at the end of the wizard, it's going to tell you exactly what the UNC name is for that share.

Advanced File Sharing
1:58-2:06
If simple file sharing is so great, why do we need advanced file sharing? Well, it gives us a lot of control over the sharing that we don't have in the basic wizard.

Custom Names
2:07-2:39
The first advantage is it's going to allow me to set a Custom name for the share; this is great, because if I add a $ to that name, I can actually hide the share. The only way to get to it is if I know it's there. All of the roots of the drives on the server will be shared out with the letter of the drive and a $. I can get to the C: drive of any computer in my network just by going to C$ for that particular computer. It only works for administrators, but I can create hidden shares and give permissions to whoever I want.

Multiple Share Names with Different Permissions
2:40-3:00
It's possible with the Advanced Sharing tab to have multiple share names with different share permissions. As far as I'm concerned, if they took that out of the operating system, I'd be just fine. I don't recommend this, because then you've got to figure out how they got into that share in order to determine what their permissions are. It's not recommended. If I use Advanced File Sharing, by default, it assigns the Everyone special identity the Allow-Read share permission.

Everyone Assigned Allow-Read Share Permission

3:01-3:22

Now, you need to be aware of that if you're going to use advanced file sharing, because you need to go in and change that. Otherwise, regardless of what NTFS rights anybody has to the share, they're all going to be locked down to read. They're probably not going to be too happy with you.

Hidden Shares and Multiple Share Names

3:23-5:13

Now let's take a quick look at Hidden shares, Multiple share names.

Suppose this is my file server, and on my file server I've got a couple of folders, Apps and Department, and inside Department I have Data. Let's say, for the sake of example, that what we use Apps for is we store the executable files for the applications in our network up here in the Apps folder. So, I'm the network administrator, I can connect up to that Apps folder from any computer in the company and use it to install whatever software we're rolling out. I don't want my users browsing through shares and say, "Hey, Apps folder, that looks like a good folder to get into." Even if I lock them out, they're going to be very interested in hacking that folder, and it's much harder to keep people out of things if they know it's out there. If they don't know it's out there, there's nothing to attract them. I'll go through, and when I share it, I'll make my share name Apps$. What that means is, when anybody looks at a list of the shares available on FS1, they would not see this name. The only way to connect to it would be to put in the UNC pathname to it: \\FS1\Apps$. It's a security measure that's really good if you have things that just should be for administrators.

Different share names gets very confusing, so suppose I took Department and I shared out as Dept, I'll give the Everyone group Read. Then I'll go through and I'll take the same folder, I'll share it out again as Dept1, and give the Everyone group Full control. This is an extreme example, but literally I have no idea what rights somebody would have accessing these folders. The information that's in there, if they connect up to the share name using Dept they're going to be able to read this data.

If they connect it to the share name user Dept1, they're going to be able to modify and delete it. This is not a great idea, it's better to share it out once with one name and set up the proper permissions.

Caching

5:14-5:29

The last benefit of Advanced File Sharing is that we're going to be able to use the Caching button to set up offline files, so that means if the users open up files from the share, it can potentially be cached on the client, so they would have access to those files when the computer's not connected to the work network.

Share Permissions

5:30-5:31

Let's take a look at our Share Permissions.

Read

5:32-5:35

Read gives me the ability to access the share. They can't open the share if they don't even have Read.

Change

5:36-6:27

Change lets me access the share and add, change, or delete content. This is my favorite permission; let me tell you why. The Change share permission is really the equivalent of the NTFS Modify permission--they both let the users do the same thing. They can add, change, and delete data. The great thing about it is, anytime I'm on an exam situation, or even in a real life situation, if somebody says a user has Modify, that tells me right off the bat that we're talking about NTFS permissions.

If somebody says that the user has Change, that tells me right away we're doing Share permissions. If they say Read or Full Control, I really don't know which set of permissions I'm dealing with unless they tell me. NTFS permissions are in play all the time; share permissions only affect me if I come across the network.

Full Control

6:28-6:54

Full Control is really the ability to change it, plus the ability to modify the share permissions.

Share permissions are cumulative, which means that the most permissive permission will apply unless I have a Deny. Deny overrides allow. If I'm a member of a group that has Change and I'm a member of another group that has Full Control, my effective permissions will be Full Control--what I can do when everything is said and done.

Best Practice for Share Permissions

6:55-7:33

Best practice for share permissions are, first of all, keep your share name short. It makes it easier for people to type, less problems getting connected up using the UNC pathname. Don't use spaces in share names. It shouldn't be a problem going through the GUI, but I actually love to connect to shares using the command prompt, and it causes all sorts of problems with command lines and scripting. Assign the Everyone group, Allow-Full Control, and lock down the content using NTFS. NTFS permissions are always in effect. Share permissions are only in effect when I come across the network, so NTFS are more secured permissions, and we're better off just sticking within that system.

Nested Shares

7:34-8:21

The last advice is do not create nested shares. Let me show you what that means. We can stick with our same example. Let's say I've shared out this Dept folder using the name Dept, and I gave Everyone Read. Then I go through and I share out Data, but here I give IT Full control. If a member of the IT department is going to use something inside of the Data folder, I really don't know what permission it will have. If that person connected to the UNC pathname using the Dept share name, then all they're going to have is Read, because that's all I've got for permissions there. If they connected using the UNC pathname with the Data in it, then if they're a member of the IT group, they'll have Full control. That's a real mess to troubleshoot.

Summary

8:22-9:25

Once you share out a folder, don't share out any folders underneath it. Our goal really is to be able to set up a system where I share out the parent folder, I set up my NTFS permissions at the parent folder, and then I walk away; let the users do whatever they want in there. I've set up my permissions correctly, and they have what they need to be able to function. We share folders in order to make the data available across the network. We can use basic file sharing, which is a wizard that will handle both NTFS permissions and share permissions, make it nice and easy, use the name of the folder as the share name, give me my UNC path, and I'm good to go.

If I need more control over the share--maybe I want to hide it, or I want to set specific permissions, or maybe I even want to set up offline folders using the Caching button--for that, I'll use Advanced Sharing. I can give out Read, which lets them open it, Change, which lets them modify the content, or Full Control, which includes Security. In practice, I'm probably going to give Everyone Full Control and lock down my data on the NTFS side, because NTFS permissions are always in effect.

6.3.2 Simple File Sharing

Simple File Sharing
0:00-0:37
In this video, we're going to take a look at simple file sharing. Let's go in and get a folder that we want to share. I'm just going to create one called Simple.
In order for users to get access to this data across the network, I have to have both my sharing and my NTFS permissions set up properly. Simple file sharing does both of those at the same time. I can go through and right-click this, Share with, Specific people. I put in who I want to share with; whatever the group is.

Specify Access
0:38-0:52
Then, I can specify their access; they can either get Read or Read/Write.
This is the same as NTFS Read and Execute. This is the same as NTFS Modify. I go ahead and hit share. It gives me the UNC pathname.

UNC pathname.
0:53-1:27
It's going to share it out with the name of the folder, so don't have any spaces in the folder; and very basic.
If I take a look in the properties of this folder, we can see what it did. We go into Advanced Sharing; it's shared out under the name Simple. If I look in Permissions it's given Everyone full control. The idea being, I'm going to lock it down on the NTFS side.
If I go in to NTFS, I can see Domain Users, and they've been given Read and Execute because I gave them Read.

Summary
1:28-1:58
Simple file sharing, very easy to use; it's going to set up both my share and my NTFS permissions which is fantastic, but it is simple. It's going to use the name of the folder for the share name. It's going to go ahead and give everyone full control on the share side. Whatever needs to be done will be done on the NTFS side and I'll be ready to connect to that share. It does not let me pick my share name. It does not do anything with offline files. Again, it's just simple file sharing. But, for a lot of the shares that you're going to create, that's fine. That's good enough.

6.3.3 Advanced File Sharing

Advanced File Sharing
0:00-0:30
In this video, we're going to take a look at Advanced File sharing. Let's go find a folder we can share. I'm going to make a new one here in the root of C:, and I'm going to call it Advanced. If I want to go through and have a lot of control over what happens when I share out this folder, I'm going to right-click it and go to Properties. I want to go over to my Sharing tab, and I want to use Advanced Sharing. I'm going to share this folder.

Name of the Share
0:31-0:38
Now, the first thing that you should see is that I can control the name of the share. It doesn't have to be Advanced. It could be ADV.

Adding a Dollar Sign ($)
0:39-0:50
If I add a dollar sign ($) in here, it'll actually create a hidden share that won't show up when people are just browsing the shares, but if I know the path to it and I type the path in specifically, I can get to that share.

Limiting the Number of Simultaneous Users
0:51-1:10
I can limit the number of simultaneous users. That's quite a lot of users, but if you do have a share where the network performance is really bad -- I very rarely have run into that, but every once in a while, I get somebody that says, you know, the files in that share are so big. We notice that if more than two people are in there, it actually cuts down on the performance of the server. Well, you certainly could do that.

Permissions
1:11-1:35
I should also go in and set up my Permissions. You should be aware that the default share permission is Everyone read. The reason you should be aware of that is, if you don't change it, that's the maximum anybody's ever going to be able to do with this share. The best practice is to just give everyone full control and lock the share down on the NTFS Permissions.
In terms of my Share Permissions, there are three permissions that I can give out.

Read
1:36-1:44
Read lets users open the share, plain access it, and open up the documents that are in the share, but they can't make any changes.

Change
1:45-1:53
Change is the equivalent of NTFS modify. It lets them create, add, delete -- pretty much everything except security.

Full Control
1:54-2:13
Full Control gives me the right to change plus modify the security for the share, and those are the only three permissions that I've got on Share Permissions.
Share Permissions only affect users coming in across the network. That's why the preference is to lock down my data using NTFS, which affects everybody all the time.

Set Up Offline Files: Caching

2:14-4:06

The last thing that I can do in this sharing dialogue box is set up off-line files, which is done in the Caching button. Kind of confusing, because they'll say offline files, but I've got to click Caching to get in there. Just be aware that caching is the same as offline files.

By default, only the files and programs that the users specify are available offline. No support for branch cache unless I turn that on. It might be that this share is very secure, so you say, well, I don't want anything to be available offline. Or maybe you've got to go all the way in the other direction and say, now, anything that they open should be available offline on their computers when they're disconnected. It really just depends on what you're trying to do to support your users.

Be very careful with offline files. It's not good to have users actively using offline files on shares that are modified both on the share and on disconnected workstations, because if they're modified in both places, it's done with timestamps, and when the computer reconnects, it'll prompt the user and say, hey, the one on the server is not the one that I had a copy of offline.

That's changed, plus the one on the workstation has changed. What do you want to do to make this work? There's no good answer to that question. I can have the copy that's on the server WIN, which will erase the changes on the laptop. I can have the laptop WIN, which will erase the changes on the server, or I can create two copies of the file, which is what I didn't want to do to begin with.

A lot of times, we're in here to turn this off, but there can be situations where it's beneficial as well. It just depends on what you're trying to achieve with that particular share.

I'm going to hit OK. Notice it's shared out as \\DC1\Adv. If I go back in, I can add an additional share name. I'm not in favor of this, because it gives me two portals into that share. I can set different permissions for the different share names, but then again, it makes it difficult to troubleshoot because I've got to look at which share name they connected to.

You can add multiple share names.

4:07-3:49

Multiple Share Names

3:50-4:51

You can add multiple share names. I'll hide this one for you, but again, I'm not in favor of this, and mostly because I don't see them both in this box. Unless I actually click the down arrow, I have no idea that it's also shared out as ADV.

If I go in and take a look at the shares on this computer, you can see I see the ADV share, but I don't see the Advanced dollar sign share. That's what it means to hide a share. I can still get in here by typing in the actual name of the share.

Summary

4:52-5:06

Those are some of the features of advanced sharing. The biggest thing to remember if you're going to use advanced sharing is that you go in and change the default share permissions. Otherwise, everything's just going to be read-only, and you won't have any shares where users can actually modify data.

6.3.4 File Share Facts

Share permissions work together with NTFS permissions to control access. Be aware of the following when managing share access:

- Share permissions are in effect only when files are accessed through the network share. If files are accessed locally, share permissions will not control access.
- NTFS permissions restrict access to both local and network users.
- Both share and NTFS permissions must be configured for a user to access the share. If a user is allowed share access, but no NTFS permissions are set for the user or a group to which the user belongs, no access is allowed.
- Share permissions are cumulative:
  - The most permissive permission will apply.
  - Deny overrides Allow permissions.
- Effective permissions to shared folders are the *more restrictive* of either share or NTFS permissions. A user's effective permissions cannot be greater than the share permissions assigned to the user or a group to which the user belongs. For this reason, a common strategy for assigning permissions is to:
  - Assign Full Control share permissions to Everyone.
  - Use NTFS permissions to control access.
  - Whenever possible, assign permissions to groups rather than users.
  - Add only necessary groups and assign only the necessary permissions.
  - Do not create nested shares.

Even though Everyone has share permissions, only the users or groups with NTFS permissions will have access.

The following table describes the two types of file sharing:

Type	Descriptions
Simple file sharing	Simple file sharing uses the name of the folder as the share name: • Folder names should be short. • Folder names should contain no spaces. • When the simple files sharing wizard completes, it provides the Universal Naming Convention (UNC) name for the share. • The UNC form is \\SERVER\SHARE.
Advanced file sharing	Advanced file sharing gives you more control over sharing than you have in the File Sharing Wizard. Advanced file sharing: • Allows a custom name for the share. Add $ to hide the share Roots of a computer will be shared out and accessible to Administrators by typing the drive letter and a $ sign (for example, C$). Administrators can create hidden shares and assign permissions to specified users.

	• Allows multiple share names with different share permissions. This is not recommended. By default, everyone has the read share permission. • Allows offline files. Use the **Caching** button to set up offline files. The default permission for folders shared using Advanced Sharing is Read.

The following table lists advanced share permissions.

Permission	Access
Read	Access share
Change	Access share; add, change, or delete content
Full control	Change, plus the ability to modify share permissions

6.4 Sharing on Server 2012

As you study this section, answer the following questions:

- Which type of file systems use the Server Message Block (SMB) protocol for file shares? When is the Network Files System (NFS) protocol used instead of the SMB protocol?
- How does access-based enumeration work with NTFS permissions to restrict access to files and folders in shared folders?
- How would you make shares available to offline users?
- What feature must you have installed to use the SMB - Share Advanced option to create a share?

After finishing this section, you should be able to complete the following task:

- Create shares on Windows Server 2012 using Server Manager.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  Share Folders and Configure Share Permissions
  - Share Folders
  - Manage Shared Folders
  - Manage Share Caching
  - Configure Share Permissions

This section covers the following 70-410 exam objective:

- 201 Configure file and share access.
  This objective may include but is not limited to:
  Create and configure shares
  Configure share permissions
  Configure offline files

Sharing on Server 2012
0:00-0:16
Let's take a look at creating shares using server manager. I'm going to click on File and Storage Services, Shares, and then we're going to go up under Tasks and create a New Share. Now we have some choices.

NFS Share
0:17-0:34
The ones that start with NFS are really used if I have the NFS file system installed. They're used to share files with UNIX-based computers. We don't actually need to go through and make any shares like that unless we have UNIX-based computers that are going to be attaching to the shares.

SMB Share
0:35-0:40
If it's Microsoft clients, then we're interested in the SMB share options.

SMB Share - Quick
0:41-1:04
Now the Quick share option is really just the fastest way to get a share up and running. It doesn't ask you for very much. I go through, I'm making a share on this particular computer. I type in my path. I've got my share name of Test. There's my path. This is going to be the UNC name of the share, hit Next. If that doesn't exist, it will prompt you to create it, and then I can select some options.

Access-Based Enumeration
1:05-1:14
Access-based enumeration means that the users will see only the files and folders that they have permission to access. If they don't have at least read permission, they're not even going to see the folder.

Caching of Share
1:15-1:22
Caching allows them to make the files available offline. If I have BranchCache installed, which I don't, I could also enable BranchCache on the share.

Encrypt Data Access
1:23-1:31
And then, encrypted access means that when people access the files across the network, in that case, it will be accessed using an encrypted channel to protect the data.

Customizing Permissions
1:32-1:50
I have the opportunity to adjust the permissions if I need to for the share. It's going to give everyone the Full Control share permission. These are just NTFS permissions that I would be customizing. I hit Next, and then I could go ahead and Create, and I have my share.

SMB Share - Advanced
1:51-1:55
Let's go back in and take a look at the Advanced option.

File Server Resource Manager (FSRM)
1:56-2:29
Now, Advanced is only going to be available if I have File Server Resource Manager installed on the computer, which in this case, I do. If you don't have FSRM installed as a feature, Advanced will give you error, and you won't be able to use Advanced to create the share.
For the first part, it's really just the same procedure. I hit Next, there's my UNC path name, going to create the folder for me, same choices that we had on the Quick sharing, same option to adjust the NTFS permissions.

Folder Usage Properties
2:30-2:38
Now I get in and I can actually select the type of data that's going to be stored in this folder. In this case, I'm going to check User Files.

Specify the Folder Owner Email
2:39-2:49
If I want to Specify a Folder Owner Email, I can do that. I just separate it out with semicolons if I have more than one. I can choose to Apply a Quota based on a template or not apply a quota, and then I'm ready to create my share.

Apply a Quota
2:50-2:58
My third option is to create an SMB Share for Applications.
SMB Share - Applications
2:59-3:48
This is going to actually disable some of the settings that we saw in Advanced so that we don't turn on anything that would conflict with certain applications like Hyper-V, SQL, other server applications.
I'm going to hit Next. I give my path; and it creates the folder for me. You can see it's turned off Access-based enumeration. It's turned off offline files. All I really select is to Encrypt the data as it's going across the network. I can adjust my NTFS permissions, and then I create my share.
That's all really there is to it to creating shares with Server Manager.
Summary
3:49-4:13
You also can certainly create shares within the file system in the way we've done it for quite some time. If you like working in Server Manager, you can go through and you can get the options right off the bat. If you create a share using Windows Explorer, you can always come in after the fact and set up AB (Access-Based) offline files, etc, inside of shares, by modifying the existing share.

6.4.2 New Share Wizard Facts

New in Windows Server 2012, the New Share Wizard allows you to create a new shared folder. The New Share Wizard is part of the Server Manager task-based interface.

The New Share Wizard allows the following protocols for file shares:

- Server Message Block (SMB) protocol is used for Windows-based file systems. When using SMB, permissions are granted to individual users and groups.
- Network Files System (NFS) protocol is used for UNIX-based file systems. When using NFS, permissions are granted to specific client computers and groups using network names.

When you create a new shared folder, you choose the profile for the share. The profile is based on the protocol and configuration settings. The following options are available in the New Share Wizard:

Option	Description
SMB Share - Quick	**SMB Share - Quick** is the fastest way to create an SMB file share. **SMB Share - Quick**: • Supports general file sharing. • Allows advanced options to be configured after the share is created.
SMB Share - Advanced	**SMB Share - Advanced** allows you to configure the following settings when you create the share: • Folder owners for access-denied assistance. When you use this option, the folder owner is notified when a user needs assistance in accessing a share. • Default classification for the data in the folder. The file classification determines the management and access policies that will be applied to the files. • Set quotas. The quota specifies the maximum share size.
SMB Share - Applications	The **SMB Share - Applications** option allows you to create an SMB file share configured for use with server applications, Hyper-V, and certain databases.
NFS Share - Quick	**NFS Share - Quick** is the fastest way to create a NFS file share: • Supports general file sharing. • Allows advanced options to be configured after the share is created.

NFS Share - Advanced	**NFS Share - Advanced** allows you to configure the following settings when you create the share: • Folder owners for access-denied assistance. When you use this option, the folder owner is notified when a user needs assistance in accessing a share. • Default classification for the data in the folder. The file classification determines the management and access policies that will be applied to the files. • Set quotas. The quota specifies the maximum folder size.

After selecting the location for the share and naming it, you can specify the additional settings as follows:

Setting	Description
Enable access-based enumeration	When selected, only the users with access to the share are able to see the share.
Enable continuous availability	Enables automatic fail over in the event the share fails.
Allow caching of share	Makes the share available to offline users. **Enable BranchCache on the file share** is a separate option for BranchCache to allow users in a branch office to access locally cached copies of a network share.
Encrypt data access	Enables encryption for remote file access to the share.

6.4.3 Work Folders

Work Folders
0:00-0:13
Let's spend some time discussing Work Folders. Work Folders are really useful. They were introduced in Window Server 2012 R2 and the neat thing about them is the fact that they allow users to sync their work data to their own mobile devices.

'Bring Your Own Device' Problems
0:14-0:58
This is a big issue in the modern working world. End users have their personal devices that they love so much that they bring them to work. In fact they bring them to work and want to use them to complete work tasks.
This presents system administrators with a key problem. Because users bring their own devices to work and use them to work on sensitive company data, this sensitive company data ends up being copied to that individual user's device and this is problematic in a lot of different ways. For example, if that device that you have no control of because the user owns it, not the organization, if that device gets lost or stolen, then that sensitive data is potentially exposed and this isn't good.

The Role of Work Folders
0:59-1:52
In order to remedy this, we can use this new feature called Work Folders. With Work Folders, users can access the work files that they need on their mobile devices as well as on their desktops and they can access these files from anywhere that they have an internet connection, but they allow you as the system administrator to maintain control over this data. We do this by storing this data on centrally managed file servers.
In addition, it allows you as the system administrator to enforce security policies on these mobile devices. For example, we may want to require that these devices have encryption enabled and we also want to make sure that there is some type of a lock screen password or PIN code assigned to these devices. Again, if this device gets lost, we want to make it as difficult as possible for someone to compromise the data on these devices.

Sync Share
1:53-2:15
With this in mind, let's take a look at how Work Folders work in a little more detail. Work Folders store user files that they're going to need to work on from their mobile devices on the server in a special type of share called a sync share. The sync share stores data that will be synced to the end users' mobile devices.

HTTPS
2:16-2:45
Work Folders will use the HTTPS protocol to communicate with users' mobile devices and sync the data to them. One of the neat things about Work Folders is the fact that you don't have to create a new share and copy data to it in order to implement Work Folders. You can actually just specify an existing folder that already contains the user data in question as a sync share. Basically this enables you to deploy Work Folders without having to do a lot of work.

Benefits of Work Folders
2:46-2:48
There are many benefits to using Work Folders.

Provides a Single Point of Access to Data

2:49-4:31

One of the key ones is the fact that we provide a single point of access to work files. A user can bring in their personal device, connect to the sync share using an HTTPS connection and access the data over here. This is really good because it eliminates the 'e-mail yourself a file' scenario where the end user here says, "Hey, I need this certain data and I want to work on it on my mobile device," so in order to get it to my mobile device, I'm going to go to my desktop system and I'm going to access the data and e-mail it over here to my mobile device and work on it. When I'm done, I'll e-mail it back to myself, to my desktop system, and I'll copy it back up to the server.

Doing this creates a huge number of problems. For example, we have version control because we end up with multiple versions of that same data, a copy on the server, a copy on the workstation up here somewhere, and a copy on the mobile device. In that scenario, I can tell you from personal experience, eventually at some point the wrong version of the data gets saved overtop of the correct version of the data and everybody has to redo a lot of work.

The other problem that it creates is the fact that we now have this data over here on this mobile device and if it's sensitive company data, this is problematic. If I go on a business trip and I forget and leave my mobile device in a taxicab and somebody else gets it, if they can get through the security of the device to the data then they've got access to that information.

By using Work Folders, we eliminate the scenario and the data is stored permanently on the server and then moved to the mobile device when it's needed but we still maintain control of the information on the mobile device--as we'll talk about in just a second.

Can Access Files While Device is Offline

4:32-5:34

Another key benefit is the fact that the user can access their files while the device is offline. If I'm an end user and I have Work Folder set up, I can use the HTTPS protocol to access the necessary data off of the central server, from our sync share. That data is pulled down to my mobile device and then I can go on the road and I can still work on my files. I don't have to have an active internet connection.

In this scenario, I pull the information down, I go on the road--say I'm sitting in a taxicab or on a bus or a train somewhere and it doesn't have Wi-Fi access so I just start doing what I need to do on my files here. Then when I have internet access again, say when I get to my hotel room, I can synchronize the changes that I made to the files here back over to the central server. This eliminates the version control problems I talked about earlier with the 'e-mail yourself a copy of the file' scenario.

The synchronization process ensures that the latest changes are always updated here on the central server. In addition, you can integrate Work Folders with your existing file server management technologies, including file classification systems, as well as folder quotas.

Compatible with File Classification and Folder Quotas

5:35-5:45

Can Apply Security Policies to Mobile Devices

5:46-6:31

And the key benefit--and I think the most important benefit--of Work Folders is the fact that you can apply security policies to the mobile devices.

Remember, we said earlier that these mobile devices are actually owned by the end user. They're not owned or managed by the organization. As such, we're relying on the end user to make sure they have the appropriate security configured. That's not a good assumption to make. End users have different degrees of technical abilities and frankly some of them just don't care. Therefore we need a centralized way to ensure that our security policies are uniformly and consistently applied to all of these mobile devices even though they're not owned by the organization; they're owned by the individual end users.

Can Encrypt Work Folders

6:32-6:35

Using Work Folders, we can encrypt the Work Folder on this mobile device and we can also require the use of a screen lock password or PIN number.

Require a Lock Screen Password or PIN

6:36-7:35

Many times, you'll find the end users for convenience will not use a PIN or password to access the device. If a device that's owned by the end user does not have a lock screen password or PIN assigned to it and the end user has copied sensitive organizational data to that device and they, say, lose that device in a taxicab-- we've got a real problem because that means anybody who finds it now has access to sensitive company information.

By using Work Folders, we can basically enforce the use of a lock screen password or PIN number on the mobile device. Unless you know the lock screen password or PIN, you're not allowed to access the device. The nice thing about it is that we don't have to rely on the end user. We can just push that down from the central server and make sure that that's configured whether the end user has the technical ability to set it up themselves or not, it's done for them.

Server Requirements

7:36-8:29

In order to use Work Folders, your server has to meet several requirements. First of all, the server has to be running Windows Server 2012 R2. This is required for hosting the sync shares that will contain the user files. The volume that we're going to create the sync share on must be formatted with the NTFS file system and you must have the Work Folders role service installed on the server. This is done by installing the file and storage services role on the system. When you do this, several key things will happen. First of all, the Work Folders page will be added to the File and Storage Services console in Server Manager, it will also install the Windows Sync Share Service which we used by Work Folders to host the sync shares, and will also install the sync share Windows PowerShell module that you can use then to manage the Work Folders on the server.

Administrative Requirements

8:30-10:05

In addition to the server requirements, there are other things that you as the system administrator have to do as well. First, you need to create a server certificate for each file server that will host a Work Folder sync share. Really, if you want to do things right, these server certificates should actually be from a trusted public certification authority, a public CA. You can use your own internal CA using a self-signed certificate, but it's going to require a little bit of work on the individual mobile devices because they're not going to trust the certificates that come from that CA, by default.

To save yourself the headache it's really easier if you go out and get a certificate from a public CA that is trusted, by default, by most mobile devices. You also need to reconfigure your firewall to make the server where the sync shares are going to reside accessible from the internet so that's going to require a little bit of infrastructure work on your part.

Finally, you need to set up a publicly registered domain name with your DNS server. I say publicly right here because if you have an internal DNS server that is not publicly registered, you're going to have a problem, because the devices are going to try to attach to the sync share on your central server using a DNS name, so you need to use a DNS name that's being replicated around the world by our root level DNS servers. If not, it's going to again require a little bit of manual configuration work on each individual mobile device and that's probably something you don't want to have to do.

Client Requirements

10:06-12:42

Now that we got the server requirements out of the way, let's take a look at the client requirements, because there are several different requirements to use Work Folders on individual mobile devices. First of all we need to talk about the operating systems that are supported by Work Folders. You're kind of limited, you have to be using Windows 8.1 or Windows RT 8.1. Notice that Windows 8 and Windows RT 8 are not listed. You got to be using 8.1. You could, however, still use Windows 7. If you decide to use Windows 7, then it has to be Windows 7 Professional, Windows 7 Ultimate, or Windows 7 Enterprise.

The home versions of Windows 7, for example, won't work. In addition, if you're going to use Windows 7, the systems have to be joined to your organization's Active Directory domain. This is not true for the Windows 8.1 versions of Windows. They don't have to be joined to the domain. In fact, with Windows RT, you cannot join the host to the domain, but if you're going to be using Windows 7, the host does have to be joined to the domain.

There has to be enough free space on the local storage for the work space. In fact that has to be an NTFS formatted storage device in order to use work spaces. It's important to note that the default work space folder will be located in the %USERPROFILE% directory in the subdirectory called Work Folders. You don't have to use that folder if you don't want to, you can define your own. By default, this is where it will be stored. The end user actually can change it themselves during the installation of the Work Folder. They can redirect it to somewhere else if they so desire.

Remember, I said earlier that it has to be NTFS formatted in order for a Work Folder to be created on that particular storage device. This becomes an issue if you're going to be using external storage devices on the mobile device for the Work Folder. For example, maybe you want to use a micro SD card or USB flash drive. If you're going to use Work Folders on that storage device, you've got to reformat it with NTFS in order for it to work with Work Folders. By default, micro SD and USB devices are usually already formatted with FAT 32 which won't work with Work Folders, so make sure that if your end users are going to be using those types of devices that they get reformatted with NTFS.

There is actually a file size limit when you're dealing with Work Folders. It's pretty generous--it's 10 gigabytes. That's a pretty big file to be synchronizing over the internet, so it probably won't be an issue, but do be aware that that limit is there and set.

Work Folder Implementation

12:43-12:47

With that in mind, let's take a look at the steps we need to complete in order to implement Work Folders.

On the Server

12:48-13:12

First, on the server, we have to complete several tasks. First, we have to install the Work Folders role then we need to create the sync share. This is done by launching the new Sync Share Wizard from within Server Manager. The sync share will map to a local path where all the user folders will be hosted. It also identifies which group of users will be granted access to the sync share.

On the Client

13:13-13:30

Then on the individual mobile devices, we need to enable Work Folders.

To do this, we go into Control Panel, then go into System and Security, and then select Work Folders. Once there, we need to specify where the Work Folders will be stored on the mobile device.

Administrative Access

13:31-14:23

Before we go any further, we do need to discuss the issue of administrative access. This is very important because by default, your administrator user will not be able to access user data stored in the sync share on the central server. Many system administrators are not going to be happy with that configuration. They want to have control over what's being stored and synchronized out of that share.

Therefore, if you want to enable administrator access to the user data, then when you are setting up the new sync share, you need to uncheck the option that's displayed that says, 'Disable inherited permission and grant users exclusive access to their files.' This option is marked by default. If you do that, then the administration will not be able to access the user data in the share. If you want to have control over the user data which I dare say just about everybody will, turn this option off.

Encryption

14:24-14:33

In addition while setting up the sync share, you can also enable encryption and you can also enable password policies that will be enforced on the user devices.

Work Folders to be Encrypted with Enterprise ID

14:34-15:11

Encryption policies will require the documents in the Work Folders on the mobile devices to be encrypted with the Enterprise ID. This is significant because data encrypted with the Enterprise ID will use a different encryption key from the EFS key that's used on the local system to encrypt the end users' personal documents on the same device. This is something it takes a little bit of getting used to because when working with EFS, we're used to having a single encryption key per user that's used to encrypt their files in the file system.

Wipe Only the Work Folders Data

15:12-16:31

When we're dealing with Work Folders, we're working with two encryption keys, one for the end user and one for the Enterprise. The reason this is significant is because it essentially allows us to separate the organization's data from the end user's personal data. This separation allows the administrator, if necessary, to wipe data from the Work Folders on the mobile device without touching the end user's personal information.

If you're talking about remote wipe and we're dealing with end users' personal devices, you can imagine they're not real keen on the idea, because if you explain what remote wipe does, they're going to be unhappy, because they'll see that the administrator can clean their device off. If they've got their own personal e-mail, their own personal pictures, their own personal files and such, their own personal music on it, they're not going to want the administrator to wipe their device.

Using encryption in this way, we protect the end user's personal data because we don't have the EFS key for the end user's personal files, the administrator can't wipe them out. The administrator can only wipe out data that was encrypted using the Enterprise ID key and that will make your end users much happier and much more comfortable with allowing you to manage the Work Folders on their personal mobile devices.

Password Policies

16:32-17:34

The password policies we talked about a minute ago enforce several different configurations on our mobile devices. We can specify a minimum password length and by default it is six characters. We can also set the screen autolock time. Again, by default, it's set to 15 minutes although you can set it to be less than that. We can also set a maximum password retry. The default value is 10, I believe, but you can set it to be less, so if the device is compromised, it's lost, if someone's trying to break into it, if they try to use the wrong password more than the set number of times, the device will lock and they're not allowed to re-authenticate.

The important thing to remember about these password policies is that if the end user's mobile device is not compliant with these three policies: password length, the screen autolock time, or maximum password re-tries, if they don't meet these minimum parameters, then the device will not be allowed to use Work Folders.

Synchronize Files Between the Mobile Client and Server

17:35-18:32

Once set up on the server and the client in order to use Work Folders, all the user has to do is simply create a document on their mobile device and then save it to the Work Folders location. The client will then sync with the server if there are any changes made in the local Work Folder. If you create a new document, it will be copied up to the Work Folder on the server. If you make any changes to that document after that, those changes will also be synchronized up to the centralized server in the sync share.

When the client initially connects to the server, the server will also notify the client if it has any changes that have been made by somebody else on the server. If the client doesn't have anything changed locally, it will then connect to the server every 10 minutes to see if any changes have been made by somebody else on the server. And you can actually manually trigger a sync action by creating or modifying a file on the mobile device under Work Folders and when you do that, that will cause the synchronization to occur.

Summary

18:33-18:58

That's it for this lesson. In this lesson we introduced you to Work Folders on Windows Server 2012 R2. We first talked about what Work Folders are. We talked about the various bring your own device problems that Work Folders address. We talked about the requirements for using Work Folders. We talked about how to set up Work Folders on the server and on the client. Then we ended this lesson by talking about how to synchronize files between the mobile client and the central server using Work Folders.

Configuring Work Folders
0:00-0:03
Now, let's take a look at how to set up work folders.
Set Up the Domain Controller or Server
0:04-0:11
I'm actually here on my domain controller--although you normally would set this up on maybe a file server-- the user's name is DC1.corp.builditrite.com, The first thing we want to do is create a group for the users that will be using work folders.
Create a Group and Add a User
0:12-0:50
I'm going to go up to Tools > Active Directory Users and Computers, and in the Tech OU I have a user named Shadow Ferrell who can log on as sfarrell@corp.builditrite.com. So we're going to go ahead and create a group, and I'm just going to call it WorkFoldersUsers. Then I'm going to add this user to that group.
Now that we've set that up, our next step is to add the Work Folders feature.
Add the Work Folders Feature
0:51-1:30
I'm going to go up to my dashboard and add a role and feature. I hit Next until I get to Roles. I need to expand the File and Storage Services role, open up File and iSCSI Services and down at the bottom we'll see Work Folders. So I'll go ahead and hit Next, Next, and then Install. It actually requires a feature that if it's not installed it would pop up and ask you add the feature which is the IIS hostable web core, but that may already be enabled on this particular machine.
Create the Sync Share
1:31-2:17
Now that my work folders have been installed, my next move is to create the sync share. We're going to go over to File and Storage Services, Work Folders. I can either click create Sync Share right here or I can go under TASKS and do a New Sync Share either way. So, you notice it says you need to allocate storage on a NTFS volume, you probably also want to create the security group which we've already done, so we're just going to hit Next.
Now we enter in a local path or I could send it up to a share that already exists. I'm going to type in C:\workfolders and if it doesn't exist it will prompt you and say, "Hey it's doesn't exist. Can I create this?" Sure.
Now I set up the structure.
Grant Sync Access to Groups
2:18-3:02
User alias means that the folders will get named with the login name. I'm actually going to choose user alias at domain, where they'll get named with the UPN. Notice I can sync only a particular folder, and in this case, I'm going to type Documents, so we will sync the Documents folder. Then I'm going to hit Next. You give the share a name, workfolders. Now I'm going to add the group that's allowed to use workfolders, so we'll hit Add, and if I type 'work', I only have one group that's named that. So the WorkFoldersUsers group gets added. They'll be the ones that have rights to create folders and work folders. Make sure if you create this, you get the user to log off and log back on again, so that they're a member of that particular group.
Set up the Device Policies
3:03-3:21
Now when I hit Next, I can set up my device policies. Normally we want to encrypt work folders. I'm not going to worry about it because this is just testing. Here's the password policy that requires the lock screen to go off within 15 minutes. A 6-character password and no more than 10 retries.
Create and Confirm the Sync Share
3:22-3:45
I'm going to leave that set and hit Next and then let it go ahead and create the sync share. So you can see that the sync share is listed and then down below the users that are listed, and I've got some free space. If I want to establish quotas I could install File Server Resource Manager and add quotas as well.
Set up the Client
3:46-4:11
Now that the server side is set up, we're going to go over to the client and set up the client. On my Windows 8 machine, I'm going to go ahead and log on as sfarrell. Now, because I'm not going to use SSL, I'm just testing, I'm going to go ahead and add the registry key that will allow me to disable it.

Add a Key to the Registry

4:12-5:16

I want to open up the command prompt as an administrator, and of course I'm going to get prompted because my regular user account is not an administrator. I'll log on as the administrator, and we want to go ahead and add a key to the registry. So I'm going to do 'reg add'. Sometimes commands are not case sensitive, but I would think the allow unsecure connection is case sensitive. REG_DWORD is a type that is probably case sensitive. As much as you can maintain the case, it's probably a good idea because the registry can be a little temperamental. In this case, it looks like it already exists. Normally it shouldn't, so we'll just go ahead and override it.

Now I've actually set it up so that we don't have to use SSL. In real life you want to bind SSL there and allow the clients to do that.

Set up the Work Folders

5:17-5:40

I should be ready to set up my work folders, so I'm going to right-click my Start button and go into Control Panel > System and Security, and go into Work Folders and I'm going to Set up Work Folders. I'm going to add my UPN. Notice it comes up with an error, and it says it can't resolve it.

Resolving the Error

5:41-8:23

They're going to be using the UPN like this. You actually have to add a record to DNS, and I wanted you to see the problem, so that if you do see this in real life you know how to correct it.

There's a couple of ways to do that. Let's go back to our domain controller. One way to fix this is to go into DNS and add a record, and that's the way I'm going to go ahead and do that. I need to open up my corp.builditrite.com and I'm actually going to make a CNAME record, although an A record would work just as well. I'm going to have the name that they're resolving be 'workfolders', and it's going to add the domain name. In this case, I can actually go through and browse for DC1 which is where they're hosted. Otherwise I can just do an A record for workfolders and put in the IP address of the server, either way. I am going to click OK. So now workfolders is going to resolve the DC1. DC1 is 192.168.0.100 so my clients should be able to locate the workfolders server. The other way to do this would be to go into Group Policy and add that URL in. It can also be done with Group Policy. In our case, DNS should work just great.

Let's go back to my client and see if life has improved. I'm going to Close my error, and just in case, I'm going to flush the DNS cache, because it probably has an answer in there that says, "Yes, there is no workfolders." So I'll do an ipconfig flushdns. I hit Next. If you get an error like this, "The connection with the server was terminated abnormally," it may still be trying to connect to the server via SSL. It shouldn't, because we added that key to registry, but it may.

The other way to get in is to go through and put in the fully qualified Domain Name of the Server. In our case, because we're not using SSL, I would put in HTTP and it's being hosted on DC1 corp folder at builditrite.com. I hit Next and it's going to connect up. And it says, all right normally the files that you save in Work Folders is stored in your PC, but now I can choose a different location, if I want to. I can put them anywhere I want. I'm just going to go ahead and leave the default. Then it asks me to accept the policy because we required a policy when we set up our sync share. Unless they choose, "I accept these policies," they're not going be able to go forward. So we'll set up Work Folders, syncing my files in the background.

Set up Work Folders

8:24-8:39

Anything that I store in Work Folders will be uploaded to the organization. And there is Work Folders. I can actually go ahead and create a new document.

Sync Now

8:40-9:18

Notice I have the ability to Sync Now. We'll synchronize it with the server. If I go into Manage Work Folders, I can see how much space I have. You've enabled quotas, they're only going to see as much as the quota is going to allow them to. They can even e-mail the tech to ask for help. If I want to stop using it, I can. If I want to manage my credentials, I can do that. In this case my password is current, so there's nothing that I need to do. That's how easy it is to set up Work Folders from the client.

Confirm Configuration on the Server

9:19-10:23

Let's just go back to our server and take a look at what happened on the Work Folders sync share. First of all, if I come into File and Storage Services, and I click on Work Folders, I can see this particular user is set up for Work Folders. If I right-click the user, I can suspend the user or I can go into Properties and I can see that, in fact, it did use the UPN at the CORP in there under workfolders and it tells me the last time the user has synced. This is good for a little bit of troubleshooting. If I go into that local path and I go into workfolders, you can see it's there. It's really just the first part of the UPN. It's not taking .builditrite.com but pretty close. All right. As an administrator I don't have permission because I didn't set that up. So work folders are great. They allow users to access data on multiple devices, even devices that are not joined to the domain without compromising corporate security.

6.4.5 Work Folder Facts

Work Folders, introduced in Windows Server 2012 R2, allows users to sync work data on all their devices. Data is stored on the organization's file servers and uses HTTPS to communicate with devices. Work Folders stores user files in a folder on the server called a *sync share*. Data stored in a Work Folder can be synced to users' mobile devices. Features of Work Folders include:

- A single point of access to work files from Windows 8.1 and Windows RT 8.1 devices.
- Automatic sync of offline files when Internet or network connectivity becomes available.
- Encrypted data transmission as well as standard file encryption on devices.
- Device management services, including device Compliance Settings, multifactor authentication, and data wipe on lost or stolen devices.
- Active Directory authentication, authorization, and group policy implementation.
- Centralized file management.

Work Folders does not support:

- Syncing of arbitrary file shares. Users can only sync to their own folder.
- Sharing of synced files or folders between users.

To use Work Folders, the following requirements must be met:

System	Requirements
Server	To use Work Folders, the server must meet the following requirements: • The server hosting sync shares must be running Windows Server 2012 R2. • The volume hosting the sync share must be formatted with the NTFS file system. • You must have a certificate for the server that will host Work Folders. The certificate should be from a CA that is trusted by client systems. For best compatibility, use a certificate from a public CA rather than a self-signed certificate from an internal CA. • The File and Storage Services role must be installed on the server. This role installs the Windows Sync Shares service. • You must reconfigure your network firewall to allow the server to be accessed from mobile devices over the Internet. • The server must have a publicly-registered domain name in order for mobile devices to resolve its IP address over the Internet.
Clients	To use Work Folders, client systems must meet the following requirements: • Clients must be running one of the following operating systems: Windows 8.1 Windows RT 8.1 Windows 7 Professional Windows 7 Ultimate

	Windows 7 Enterprise Windows 7 devices must be joined to your organization's domain to use Work Folders. However, Windows 8.1 hosts do not have to be joined to the domain. • There must be enough free space on a local NTFS volume to store files from the sync share on the server. Work Folders stores data in **%USERPROFILE%\Work Folders** by default. This location can be modified during setup.

To implement Work Folders, do the following:

- On the server:
  - Install the Work Folders role.
  - Use the New Sync Share Wizard in Server Manager to create the sync share.
  - Specify the users who are allowed to access the sync share.
  - Use DNS Manager to add a new CNAME record named **workfolders** that points to the A record for the server where Work Folders has been installed.
- On the client:
  - Go to Control Panel and access **System and Security** > **Work Folders**.
  - Specify where to store Work Folders on the device.

6.5 Effective Permissions

As you study this section, answer the following questions:

- What is the role of the *access control list (ACL)*?
- What strategy can you use to combine and manage NTFS and share permissions?
- What are the sources of permissions that are added together to create effective permissions?

After finishing this section, you should be able to complete the following task:

- Configure combined NTFS and share permissions.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  Manage Combined NTFS and Share Permissions

This section covers the following 70-410 exam objective:

- 201 Configure file and share access.
  This objective may include but is not limited to:
  Configure share permissions
  Configure NTFS permissions

6.5.1 Effective Permissions

Effective Permissions
0:00-0:32
Let's talk about Effective permissions, and make sure before you watch this video that you have a good grasp of both NTFS and Share permissions, so that you'll be able to follow along with the discussion.
We're going to take a look at some examples, and I'll show you how to figure out the Effective permissions. Both Share and NTFS permissions work very similarly. Within the permission sets, they're cumulative, except for Deny, which overrides Allow permissions. So if I belong to multiple groups, I'm going to get the most permissive permission unless I've been denied.

Share and NTFS Permissions
0:33-1:01
What we need to take a look at is what happens when both Share and NTFS permissions have been set up. How do we know what the Effective permissions are going to be if somebody accesses that file over the network? If they come in and they're not coming in over the network--they just sit down-- only NTFS permissions are in play. So Share permissions affect people coming across the network. We're looking at situations where both of these sets of permissions would be in play, and we want to find out what the user can actually do, based on the security settings.

Example 1
1:02-1:23
So I've set up an example here. I've got a file server 1 (FS1), which has a folder in it called Data, and that's been shared out. This is my Active Directory Structure, so IT is in OU. In there, I've made myself a little Shad account. And this user, Shad, is a member of IT Global Group and EMP, like an Employee's Global Group. So there's my user account, and I'm a member of each of these groups.

Share and NTFS Permissions Grid
1:24-2:23
Now, when I have to deal with both Share and NTFS permissions, I just make a little grid exactly like this; Share on one side, NTFS on the other. So from a Share perspective, the Everyone Special Identity has been given the Read permission. The IT Global Group has been given the Change permission. On the NTFS side, the user Shad has been given Read, IT Global Group has been given Full Control, the Employee Global Group has been given Read. And now I want to know, what can this user Shad actually do when he connects up to this share.
The way you approach this is to look at each set of permissions individually. So if we take a look at our Share permissions, everyone is a member of the Everyone group, so that gives me Read, but since I'm a member of the IT Global Group, I also have Change. So, my Effective Share permissions would be Change. Just looking at NTFS, I get Read, because that's been given to me personally. I'm a member of IT Global, which gets me Full Control, and I'm also a member of Employee Global, which gets me Read. So, my Effective NTFS permissions would be Full Control.

Cumulative
2:24-3:06
Within the permission sets, you get the most permissive permission, so I like to think of it as the cumulative. I accumulate permissions and I have everything, unless there's a Deny. If you have both permission sets in play, you total up your Share permissions to find the Effective Share permissions, total up you NTFS permissions to find the Effective NTFS permissions, and then you get the lesser of the two cumulatives. So my overall Effective permissions here would be Change.
Pretty much every textbook you will ever read about NTFS permissions and Share permissions use the terms most permissive, least permissive or most restrictive, least restrictive. That's why I try to use a little bit different language. So, total up the Effective permissions for each permission set, and then you get the lesser of the two totals.

Example 2 (Deny)
3:07-3:53
Let's take a look at another example. We have our same share, same users, same groups, but now the permissions have changed a little bit. Again, we just look at each set independently. So on my share side, I'm a member of Everyone, and that gets me Read. I'm a member of IT global, and that gets me Change. So my Effective permissions are Change.
On the NTFS side, my particular user account has been denied Read. I'm a member of IT Global, which gets me Full Control, and I'm a member of Employees Global, which gets me Read. Well, permissions are cumulative, except for Deny. So this Deny-Read is going to override all of this, and that's going to be my Effective NTFS permission. Once I've totaled up each side, I get the lesser of the two totals. So my overall Effective permissions in this scenario is going to be Deny-Read.
We'll look at one last example.

Example 3 (Best Practice)

3:54-4:22

Here, on the Share side, they've actually followed best practice. Everyone has been granted Full Control. We're not sure why, but they went ahead and granted IT Global Change as well. So my Effective Share permissions are going to be Full Control. On the NTFS side I've been granted Read. IT Global has Full Control, Employee Global has Read, but my cumulative Effective permissions are Full Control. Here it doesn't matter which side I circle, because the permissions are the same, but hopefully you can see why Everyone Full Control is not a problem.

Security

4:23-5:05

NTFS permissions affect everyone all the time. Share permissions only affect people coming in across the network. In an ideal production environment, users will only be accessing the server across the network. That being said, there's no such thing as being too cautious or too secure. You really want to err on the side of security and caution as a network administrator. So best practice is, open it wide up on the share side. You can safely give everyone Full Control as long as you lock your data down on the NTFS side, because if NTFS comes up less than Full Control, they're going to get the lesser of the two totals. If NTFS comes up Full Control, that's the only situation in which the users would actually have Full Control rights to that particular data.

Summary

5:06-5:17

Again, when both permission sets are in play, our Effective permissions are the cumulative of all the groups I belong to, except for Deny, which overrides. I total up each set of permissions, and then what I actually get is the lesser of the two totals.

6.5.2 Effective Permissions Facts

Users and groups are added to the *access control list* (ACL) of a folder or file. The ACL entry identifies the actions that can be performed. In many cases, however, a user will have more or fewer permissions than what might be shown on the ACL. As you try to determine the permissions any one user has to a folder or file, it's important to identify the effective permissions. Effective permissions are the sum of all permissions from the following sources:

Source	Description
Explicit Assignment	An explicit assignment exists when an object is added to the access control list (ACL) of a folder or file. The ACL entry identifies: • The user or group with permission. • The specific permissions assigned to the user or group. • Whether the permissions are allowed or denied.
Group Membership	All users who are members of a group have the same permissions that are assigned to the group. When granting NTFS permissions, best practice is to assign permissions to groups. Users then obtain permissions through group membership, instead of permissions assigned directly to users.
Inheritance	*Inheritance* means that permissions granted to a parent container object flow down to child objects within the container. Inheritance as a general principal works with most types of security assignments. For example, a user given the Read permission to a folder has the Read permission to all files within the folder.

When determining a user's effective permissions, remember that NTFS permissions are cumulative. To find a user's effective permissions, examine the access control list of the target file or folder. Look for:

- Permissions the user has for the object, including inherited permissions.
- Permissions for every group the user belongs to, including inherited permissions.
- The Allow or Deny settings for each permission.
  Deny permissions *always* override Allow permissions. For example, if a user belongs to two groups, and a specific permission is allowed for one group and denied for the other, the permission is denied.
  Explicit permissions override inherited permissions, even Deny permissions. If an object has an explicit Allow permission entry, inherited Deny permissions does not prevent access to the object.

You can also use the Effective Permissions tab to view the effective permissions of any user to a folder or a file.

Effective permissions to shared folders are the *more restrictive* of the share or NTFS permissions.

The following table lists several scenarios of combined share and NTFS permissions. In this scenario, the D:\Reports folder has been shared as the Reports shared folder. The user Mary is a member of the Sales group as well as Everyone and the Users group. For these examples, assume that no other permissions exist except for those listed.

Scenario	Reports Share Permissions	D:\Reports NTFS Permissions	Mary's Effective Permissions
Scenario #1	Everyone = Change	Users group = List contents, Read & Execute, Read	Mary's effective permissions = List contents, Read & Execute, Read Mary gains share access because she is a member of Everyone. The NTFS permissions are more restrictive than the share permissions, so the NTFS permissions are her effective permissions.
Scenario #2	Everyone = Read	Sales = Modify, List contents, Read & Execute, Read, Write	Mary's effective permissions = List contents, Read & Execute, Read Mary gains share access because she is a member of Everyone. The share permissions are more restrictive than the NTFS permissions, so the share permissions are her effective permissions.
Scenario #3	Administrators = Full Control	Sales group = List contents, Read & Execute, Read, Write	Mary's effective permissions = None Mary is not a member of the Administrators group, so she does not get access to the share at all. Mary does get local access to files when logged in, but does not have network access.
Scenario #4	Everyone = Full Control	Administrators = Full Control	Mary's effective permissions = None Everyone has access to the share, but there are no NTFS permissions granted for

			Everyone or any group that Mary belongs to. The NTFS permissions are more restrictive (no access), so they take effect.
Scenario #5	Everyone = Full Control	Everyone =List contents, Read & Execute, Read, Deny Write Sales = Write, List contents, Read & Execute, Read, Write	Mary's effective permissions = Modify, List contents, Read & Execute, Read NTFS permissions for the combination of Everyone and the Sales group are more restrictive and therefore take effect. The Deny permission overrides any allowed permission from other groups or the share permissions.

Changed permissions might not take effect for currently-logged on users. Users might need to log off and back on again to get the updated permissions.

Use these suggestions to help you plan NTFS permissions.

- Identify the users and their access needs based on the actions they need to be able to perform.
- Create groups for multiple users with similar needs, and then make users members of groups.
- Assign each group the permissions appropriate to the group's data access needs. Grant only the permissions that are necessary.
- Consider inheritance when assigning permissions. Set permissions as high as possible on the parent container and allow each child container to inherit the permissions.
- Override inheritance on a case by case basis when necessary.
- Use the Deny permission carefully.

7.1 Print Servers

As you study this section, answer the following questions:

- What is the difference between a *print device* and a *printer*?
- When would you add the LPD service when configuring the Print Services role?
- Under which physical printing configurations would you choose a local printer when adding a printer? When would you choose a network printer?
- Under which circumstances would you configure a printer to use multiple print devices? When would you configure multiple printers for a single print device?
- What services are included in the Print and Document Services role?
- How would you ensure important print jobs will automatically be printed before any other print jobs?
- How can you use printer permissions to designate a user as a print job administrator?

After finishing this section, you should be able to complete the following tasks:

- Configure printer pooling.
- Restrict printer access.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  Manage Printing
  - Create, Share and Manage a Printer
  - Configure Printer Pooling
  - Restrict Printer Access

This section covers the following 70-410 exam objective:

- 202 Configure print and document services.
  This objective may include but is not limited to:
  Configure drivers
  Configure printer pooling
  Configure print priorities
  Configure printer permissions

7.1.1 Printers and the Printer Server Role

Printers and the Print Server Role
0:00-0:12
All right, let's talk about print servers and printers. Basically, what I want to do is make sure you have a good handle on the vocabulary before we go in and talk about how these things actually work. Believe it or not, the vocabulary is a little bit strange.

Print Server Role
0:13-0:22
First of all, in order to make your server into a print server, you need to install the print server role. That's going to bring in some role services, and there are a few of them.

Internet Printing
0:23-0:27
One of them is internet printing, which allows you to manage your printers and send print jobs over the internet.

LPD Printer Service
0:28-0:36
Another one is the LPD printer service. What that does is allow UNIX clients to use a Microsoft print server.

Document and Scan Services
0:37-0:43
The third role service, besides just plain printing, is document and scan services, that lets you receive scanned documents.

Definitions: Print Server, Print Device and Printer
0:44-1:42
In terms of vocabulary, the important thing to remember is the print server is whatever computer sends the print job to the physical device. Here's where the vocabulary gets tricky.
Normally, in real life, if somebody said, "Shad, what would you call a hardware device that puts ink on paper and spits it out?" I'd say, "A printer." But, it's very difficult to distinguish between the hardware and the software. We call the hardware the print device. We call the software--the icon inside of Windows--the printer. If I say, "I'm going to go ahead and install a printer," I mean I'm installing the driver. I'm creating that icon that's going to have the print queue, I'm talking about the software. If I'm talking about the hardware, I tend to use the term Print Device, so that if we're just having a conversation, you know which one I'm talking about. Then, print server is whichever computer has the printer on it, and that's the printer that sends the print job to the print device. Keep in mind, print server actually sends it to the device. If it's the hardware, then it's a print device. If it's the software, it's a printer.

7.1.2 Adding a Local Printer

Adding a Local Printer
0:00-0:03
In this video, we're going to take a look at how to add a local printer. The first thing that I need to do is add in my Print Server role.
Add Print Server Role
0:04-0:29
So we're going to Add roles and features, and notice it's Print and Document Services. I'm going to pull in some tools-- great--and there are four role services with this particular feature. By default it just checks Prints Server.
Internet Printing and LPD
0:30-0:55
Internet printing allows me to send jobs and manage the queue over Web pages. LPD provides support for UNIX clients to print to Microsoft printers. Technically, because UNIX clients can access the Internet, both Internet printing and LPD would support UNIX clients, but LPD is just for UNIX clients. Internet printing is to allow anybody to come in and manage print jobs using the Internet.
Distributed Scan Server
0:56-0:59
Distributive Scan Server allows you to receive scanned documents.
Optical Character Recognition: OCR
1:00-1:34
If you're having trouble with scanned documents having the text recognized-- maybe it's all pictures, it's not turning them into text--that technology is called OCR, or Optical Character Recognition. There's a feature--we'll go back to Features--this Windows TIFF IFilter can help out with optical character recognition; so if you're having trouble with that, you'd want to turn on this feature.
We're just going to install the print server, and I just want to make you aware of what the other role services actually do. Once I've got my print server installed, then I can go ahead and add printers.
Adding Printers
1:35-2:54
I'm going to go up under Tools, and I want to open up Print Management. Here's my Print Server--I could add other Print Servers in if I have more in my environment--and I'm going to go down to Printers. We're going to right click Printers and Add a Printer. Now, right off the bat, it gives me the choice to add it by IP address, a particular local port, or I could actually create a new port and a new printer. You would use this one if the print device is already turned on and plugged in somewhere. A lot of print devices say, "Install the driver before you turn on the device." If it's a network attached printer, meaning it's attached by a network cable and it's got an IP address, but I want to actually install it on the server before I turn the physical device on, you could come down here and Create a new port, Local Port, and then you'd put in the IP address of that print device, whatever it's going to be, and it would create that port. We don't have a network attached print device, so I'm going to pretend that it's attached to LPT1. There is probably no printer left today that would be attached to LPT1--that's the old printer port-- but we're just going to use it, because we don't have a physical device. So I'm going to hit Next.
I could use an existing driver if it's the same type of print device as I've already installed here on the print server.
Printer Driver
2:55-3:22
I don't have any print devices on this print server, so I'm going to Install a new driver, and notice I can get drivers from the Windows Update site. I could also click Have Disk if my particular print device does not show up in the list. I'm randomly picking a print driver. Now I give the printer a name.
Name the Printer
3:23-3:28
In the wizard here, I can opt to share it.
Share Name
3:29-3:40
If you're going to share it, give it a better share name than whatever it's picking here. Ideally, a share name should not have any spaces in it. If you want people to be able to search Active Directory by location, make sure you fill out the Location box.

Location
3:41-4:12
You may even want to add a Comment about where this print device is physically located. The more documentation you can do up in Active Directory, the better. We're all set. I have the opportunity to print a test page directly from here. I can check and stay in the wizard and add another printer, but I'm just going to go ahead and finish. That's how we install the print server role and install a local printer into the print server.

7.1.3 Printer Properties

Printer Properties
0:00-0:02
We're going to talk about the properties of the printer.
Print Priorities
0:03-2:02
The first thing I want to show you is Printer Priorities.
Here's my print server, PS1, and here's the printer, or the icon on that server. That icon has been shared out, and it sends jobs over to the print device, going from the software to the hardware. What's happening is the boss comes in and says to you, "Look, that guy Shad, he prints 100 documents a day, and I very often find out that I'm waiting, because there are 50 Shad documents in this queue, and there's one boss document below that, and all those first 50 documents have to get processed by the print device before the boss's print job prints." He says, "You need to make a change. You need to make sure that if I print something, my documents are going to jump ahead of anybody else's." For that, we use Priorities.
In Printer Priorities, I need two printers that point to the same print device. If you can remember nothing else about Priorities, that's key. Two printers, one print device.
I'm going to go ahead and make another printer, and maybe the original printer was called Employees, or something like that. I'm going to name this one Boss, it has a different name. It's going to point to the same print device. The difference is this: the boss's printer needs to have a higher priority than the employees'. The priorities range from 1 to 99, so I'll go ahead and just give my employee printer a priority of 1; I'll give the boss the priority of 99.
The higher the number, the higher the priority. What's going to happen is this: all of the regular employees' will print to the employee's printer. Their jobs will get sent to the print device. If the boss prints to the boss's printer, the employee printer will finish the print job it's working on. It's a common question most people ask--what happens if it's in the middle of a job? It doesn't just cut it off at page three and start on the boss. It will finish that print job, but then it will put that employee printer queue on hold, process any jobs in the boss print queue, and once all the boss's jobs have been printed, it'll go back and start working on the jobs in the employee queue. Remember print priorities--two printers, one print device. The higher the number, the higher the priority.
Printer Pooling
2:03-3:04
Let's take a look at printer pooling. I had my print server here. I've got my printer, and I've got my print device. This printer is sending jobs to this print device. In this scenario, my poor little print device is so overworked that it can't handle all the print jobs. People are waiting for hours for their print jobs to come out. What I'd like to do is have multiple print devices that can service these print jobs.
I go out and I buy some more print devices, and the key here is to make sure I buy print devices that can use the same driver as the existing print device. If I have a laser jet printer, I can buy another laser jet printer. As long as they can all use the same driver, I'm going to be good to go. Let me buy a couple more print devices.
I've got three different physical print devices. What I don't want to do is make three different printers on the server, because then we've got people guessing, "I wonder which one is busy right now, let me try printing to A, let me try printing to B, let me try printing to C." What I'm going to do is Enable printer pooling. In printer pooling, I have one printer, multiple print devices.
Ports Tab
3:05-3:29
Printer pooling is done on the Ports tab. Normally speaking, I can pick one port for that printer to send the job out to, whether it's an IP address, or a USB, it doesn't matter how the print device is connected to the print server. When I check Enable printer pooling, it's going to allow me to check multiple ports. Let's say my three print devices each have a different IP address. I can have that one printer send jobs out to all three of those IP addresses.
Port Monitor
3:30-3:58
There's actually a little piece of software called the Port Monitor that checks in with the hardware devices to see which one is free. Whichever one is free at the time the job comes in, it will send it to that particular print device. One thing to note here is the users don't know which print device is actually going to service their job. You want to make sure you locate all of these print devices in the same general area. It is not a good idea to locate them at either end of a very long hall, and watch the users running up and down. That is not best practice.

Permissions
3:59-4:07
Let's finish up by taking a look at permissions. These are the permissions that we can assign with printers. They're not cumulative, they're just separate permissions. The first permission is pretty obvious.
Print Permission
4:08-4:13
The print permission lets me print. By default, it also lets me manage my own documents. If I printed it, I can delete it.
CREATOR OWNER
4:14-4:46
The reason for that is that the CREATOR OWNER special identity, by default, has been given the Managed Documents right. If somebody's asking you, "I want to give Shad the ability to print and manage his own documents," there might potentially be two correct solutions to that.
Maybe you're just going to say, "Give him print," and that will work. If that's not the right answer, and it's not working, maybe somebody has come around and messed with this CREATOR OWNER, and taken that off. In that case, you would have to go and put this back in. This is not a security problem at all, it just means if I created the print job, I can manage it, which means I can delete it.
Manage Documents
4:47-4:57
Manage Documents, I tend to think of it as managing other people's documents. It lets me manage any of the documents in the print queue, so I can delete them, I can change priorities-- whatever I need to do with those documents, I can go ahead and do that.
Manage Printer
4:58-5:03
The last right is Manage Printer, which allows me to change settings on the printer, on the icon itself.
Summary
5:04-5:44
Keep in mind some of these special properties. I could set up the priorities, where I have multiple printers, one print device, the idea being that users who print to the printer with the higher priority will have their jobs processed first. If I have a great volume of printing, I can do printer pooling, where I've got one printer, multiple print devices, making sure all those print devices can share the same driver. That will allow it to service any items in the queue by printing it on whichever print device is available. Finally, permissions: Print, which lets me print and manage my own documents, Manage Documents, which lets me manage other people's documents or all the documents in the queue, and Manage Printer, which lets me change the settings of the printer itself.

7.1.4 Configuring Printer Properties

Configuring Printer Properties
0:00-0:24
In this video, we're going to take a look at configuring printer properties. To manage my printer, I want to go into Print Management. I'm going to go up under Tools, Print Management, and here's my printer.
I want to right-click my printer and go into Properties, and here are the properties of the printer. We'll just take it tab by tab.

General Tab
0:25-0:37
First of all, if you want people to be able to search by location, you should fill out the location box. There's also a button in here to print a test page. I want to. That's pretty much all we use this tab for.

Sharing Tab
0:38-1:21
On the Sharing tab I can Share the printer, if I didn't do that in the wizard. I also can say that I want to Render the print jobs on the client computers, or if I uncheck it, it will render the print on the server. Render means to translate into the language of the printer. The reason we want to render the jobs and the clients is that way the print server has less work. If for some reason that were causing problems, we could actually do it up on the server, too.
It's just going to put more stress on the print server that we don't need to put there unless there's a reason. If I check List in the directory, it will actually list this particular printer in the Active Directory. You can see that my Additional Drivers button is grayed out.

Additional Drivers
1:22-3:09
This you want to really be aware of. Here are the politics: with Windows Server 2012, there are only 64-bit versions of Windows Server; however, you may still have 32-bit clients in your environment.
If the print driver supports 32-bit clients, the Additional Drivers button would be colored in, and I would be able to go in and enable the 32-bit driver so that the clients can pull it down from the print server. When they install the shared printer, they get their driver from the print server. There are different types of drivers: some drivers support both 64 and 32-bit clients; some just do 64-bit. This one just does 64-bit.
If that's the case, there's really two things to do. One would be to try and get another driver from the vendor that supports both. If I have a driver in there that supports both, I can just click additional drivers and check x86 inside the dialog box. If the vendor has separate drivers for 64-bit and 32-bit, that's a different ball game.
I would need to go into the Properties of my print server, go over to Drivers, and add the 32-bit driver in here. You can see right now there's just one driver that supports 32-bit, which is this enhanced Point and Print compatibility driver, which doesn't apply to any particular device. Let's go back into the properties of our printer. That's our sharing tab--that's pretty much of all the options in there.

Ports Tab
3:10-3:15
The Ports tab lists the port that the printer is using, if I have multiple print devices that I want to have use the same printer.

Enable Printer Pooling
3:16-4:00
One icon, multiple physical print devices. I can come in here and Enable printer pooling. Notice if I choose a different port, the check mark is hopping around. When I enable printer pooling, I can select more than one port, because the assumption is I have multiple print devices. Make sure those print devices can all use the same driver, because they're going to be using this driver, and put them in the same location, because we have no control over which print device is actually going to service any particular job.
If this printer went down, the actual print device failed, and I have another print device that can use this driver.

Add Port
4:01-4:57
It's got a different icon, let's say it's salesprinter2. I can actually redirect all the print jobs from this printer to the other printer. Basically what we do here is add a port--a local port--but then I would put the share name of the other printer. If I hit ok, that becomes the port for this printer. Anything that flows into salesprinter1 will go out this shared port over to the icon for salesprinter2, and then out to the salesprinter2 device.
This is a way to kind of quickly direct an icon or printer for a print device that's failed without having to reconnect all the users to salesprinter2. On my advanced tab, I can set up when the printer's available; so I can say it's only available from certain hours.

Advanced Tab
4:58-4:57
Availability
4:58-5:04
Print Priorities
5:05-5:30
I can also set up my print priorities. If we want to have certain users have priority over other users, we would create two printers, two icons, and give them different priorities; the higher the number, the higher the priority.
I can create salesprinter2, give that a 99 priority. The print device will service jobs in the salesprinter2 queue before it takes jobs in this queue. I can also add a driver in here, set up spooling--this we don't usually mess with.
Add a Driver
5:31-5:30
Print Spooling
5:31-6:00
It should actually use the spooler. The print spooler is what allows you to get back into the document immediately. So you click Print sometimes at the bottom you'll see it's printing, but you can go back and start working on your document right away.
If we don't spool and we print directly to the printer, the user would have to sit there and wait till the job is printed before they actually get control of the computer back.
Hold Mismatched Documents
6:01-6:38
Out of everything in here, one of my favorite checkboxes is Hold mismatched documents. If you've ever had a situation where you've got a printer; it's loaded up with letter-size paper, but somebody comes in and prints a legal-size document. And then the printer hangs out with the light blinking because it's waiting for somebody to put legal-size paper in so it can service that job before it can move on to the next one.
If I check "Hold mismatched documents" because that legal job is mismatched to the paper that's actually in the print device, it will put that job on hold and let the other letter-sized jobs move around it in the queue. Color Management is specific to this printer, so there'll always be a tab that's specific to that particular print device.
Printer Specific Tabs
6:39-6:51
It's not part of Windows.
Same thing with Device Settings-- specific to that print device, but Security is on all printers.
Security
6:52-6:53
Print
6:54-7:30
Print gives me the ability to print to the printer. Now, because CREATOR OWNER has been given the Manage Documents permission, anybody who has the print right can also manage their own documents.
If I print to this printer, I can go in and delete my print job from the queue. If you're looking to allow people to manage their own print jobs, as long as CREATOR OWNER is in here you can just give them print. If somebody had pulled this off, you would need to add it back in--CREATOR OWNER Manage Documents.
There's no problem with that; it just means if I printed it, then I can delete it.
Manage Documents
7:31-7:36
"Manage documents" means manage documents in the queue so I can delete them. I can change the priority of them and then "Manage this printer" lets me adjust the Properties inside the printer.
Manage This Printer
7:37-7:48
This whole dialog box that I'm in--I need the "Manage this printer" right in order to be able to make any adjustments in here.
Those are the properties of the printer.

7.1.5 Print Server Facts

The following table lists key print and document services definitions with which you should be familiar.

Term	Definition
Print server	The computer that sends the print job to the physical device.
Printer	The software inside the print server that can be configured to send output to a print device.
Print device	The physical device connected to the print server where print output occurs.
Print driver	The software that allows the printer to communicate with the print device.
Print queue	The portion of the hard drive where print jobs are stored before going to the print device. By default, print spool files are stored in the system drive in \Windows\System32\Spool\Printers. For best performance, move this to a separate drive.
Printer port	The means by which a print device connects to a print server (parallel port, USB, or NIC).

To install a printer in Windows Server 2012, you must first install the Print and Document Services role. The Print and Document Services role is composed of four services:

Service	Description
Print Server	Adding the Print Server service installs the Print Management snap-in, which is used to: • Manage multiple printers or print servers • Migrate printers to and from other Windows print servers You can still add, share, and manage printers through the Control Panel without adding the Print Services role.

LPD Service	The LPD (Line Printer Daemon) Service uses the LPDSVC service (TCP/IP Print Server) to allow systems using the Line Printer Remote (LPR) service (typically UNIX- and Linux-based machines) to print to shared printers. • Adding the LPD Service configures an inbound exception in the Windows Firewall with Advanced Security for port 515. • The LPDSVC service does not automatically restart when you restart the Print Spooler service. • You can print from Windows to a UNIX print server if you install the LPR Port Monitor feature. LPD and LPR Services are deprecated starting with Windows Server 2012.
Internet Printing	The Internet Printing service allows you to use a Web site to print to, share, and manage printers through a Web browser. • To view a list of Internet-enabled printers, go to: **http://*print_server_name*/printers.** • To print, computers must have the Internet Printing Client installed. On Windows Vista, Windows 7, or Windows 8, use Control Panel to turn on the Internet Printing Client feature. On Windows Server 2008, 2008 R2, or Windows 2012, add the Internet Printing Client feature using Server Manager.
Distributed Scan Server	Adding the Distributed Scan Server service installs the Scan Management snap-in, which you can use to manage network scanners and configure scan processes. This service allows you to receive scanned documents from network scanners and route them to the correct destinations. The Windows TIFF IFilter performs Optical Character Recognition (OCR) and can improve the processing of scanned text.

The Fax Server role allows administrators to monitor and manage multiple fax machines remotely. The Fax Service Manager allows you to automatically make fax connections available to users and computers.

The following table identifies print device properties useful to control print jobs.

Tab	Description
General	Properties on the **General** tab include: • **Location** allows you to identify the location of a printer that users will recognize. • **Print Test Page** allows you to print a page to ensure the printer is set up correctly.

Sharing	If you did not set up the printer to be shared, you can specify sharing on the **Sharing** tab. • The **Share name** should identify the printer and not contain blank spaces. • **Render print jobs on client computers** reduces resources requirements on the print server by having the client translate the print job into the language of the printer. • To have the printer listed in Active Directory, check **List in the directory**. • **Additional Drivers...** allows you to add drivers for clients that may not have a driver for installed for the printer. There are only 64-bit versions of Windows Server 2012. If the print driver supports both 32-bit and 64-bit clients, the **Additional Drivers...** button is available and you can enable a 32-bit driver. The clients can download the driver from the printer server. If the print driver supports only 64-bit clients, you must add a driver to support the 32-bit clients. ▪ From the print vendor, obtain a driver that supports both 32-bit and 64-bit clients. ▪ If the vendor has separate drivers for 32-bit and 64-bit clients, go into print server properties and select the **Driver tab**. Click the **Add...** button and install a 32-bit driver.
Ports	The **Ports** tab lists the ports available to assign to the print device. Printer *pooling* uses a single printer object to represent multiple print devices. To pool printers: • Check a port to assign to the print device. • Enable multiple print devices to use the same printer. With printer pooling, users send print jobs to a single printer. The print server decides which print device to send the job to, performing load balancing between the printers. When creating a printer pool, all print devices in the pool must use the same printer driver. Because users can't control which print device is used for the print job, pooled print devices should be in the same physical location. • Redirect print traffic from a failed print device to another print device by adding a port for the failed print device: Add a local port with the share name of the print device to which traffic is being directed. Both print devices must use the same print driver.
Advanced	The **Advanced** tab options allow you to:

	• Select **Available from** to specify the time that the printer is available. Otherwise, you can select **Always available.** • Specify the **Priority**. A higher number indicates a higher priority. To set a different print priority for different group: Set the print priority for the first group. Add a second logical print device and set the priority for the second group. • Select **Start printing immediately** to start printing before the document is completely spooled. This option releases control of the document to users sooner than **Start printing after last page is spooled**. • Select **Print directly to the printer** only for a non-shared printer or when a program has its own spooling process. • Choose **Hold mismatched documents** to avoid delays caused by the printer waiting for the correct size paper for the print job to be inserted in the printer.
Security	The **Security** tab allows you to specify permissions for the print device: • The **Print** permission allows users to send print jobs and manage their own documents in the queue. • The **Manage Printer** permission allows users to change printer configuration settings and permissions. • The **Manage Documents** permission allows users to manage all documents in the print queue, such as pausing, reordering, or deleting print jobs. • The **Creator owner** permission grants the creator of a print job the manage documents permission. The owner can pause, reorder, or delete print jobs.

7.2 Print Management

As you study this section, answer the following questions:

- How do clients obtain the correct driver for shared printers?
- How can you set up a filter to notify you when a printer is low on toner?
- When would you choose to isolate a print driver?
- What is the advantage of location-aware printing?
- What is the difference between listing a printer in Active Directory and deploying a printer with Group Policy?

After finishing this section, you should be able to complete the following tasks:

- Configure printer pooling.
- Manage printer drivers.
- Configure printer filters.
- Restrict print access.
- Use Group Policy to deploy printers to computer and user accounts.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 5.0 File and Print Services.
  Manage Printing
  - Configure Printer Pooling
  - Restrict Printer Access
  - Deploy Printers with Group Policy

This section covers the following 70-410 exam objective:

- 202 Configure print and document services.
  This objective may include but is not limited to:
  Configure the Easy Print print driver
  Configure Enterprise Print Management
  Configure drivers
  Configure printer pooling
  Configure print priorities
  Configure printer permissions

7.2.1 Print Management

Print Management
0:00-0:17
Let's talk about Print Management and the Easy Print driver. Print Management is a utility for managing your print server that I think is really cool. What's neat about it is not only do we have it in Windows server 2012, 2008, 2008 R2, but starting with Windows 7, it's also in our client operating system.

Features
0:18-0:25
Let me tell you some of the great features about this software. One of the things I can do is set up filters that will watch for conditions on my printers.

Filter
0:26-1:15
I can set a filter that will catch any print devices that are low on toner or low on ink, or go into the Not Ready status. I can also set up a notification on the filters, so that if any of the printers enter those status and drop into the filter, I could even get an e-mail sent right to my e-mail. This particular print device is low on toner, I get an e-mail, or better yet, maybe I'll send the user an e-mail.
I also can go in and use it as a central point for managing all of my print servers. I can add all my print servers in there. I can also centrally manage my drivers and the printers themselves. So we'll talk when we get into the software about drivers and providing support for 32-bit clients, because Windows server 2012 only comes out in a 64-bit operating system. There's a couple of different things that you can do if you still have older clients.

Driver Isolation
1:16-1:44
Another really cool feature of print management is something called driver isolation. Printers are a nightmare to support, to be quite honest. You'll spend a lot of time in your career troubleshooting printers. There's all different kind of drivers: some drivers come with the operating system, some come with the DVDs or CDs that come with the print device, sometimes you download them from the manufacturer's Web site, and they're all different. Some computers might like one driver, another computer might like another driver, and that's what I see in the field, in terms of practical experience.

Print Spooler
1:45-2:57
If you have print drivers that are flaky--traditionally a common problem that we see with Microsoft print servers and client printers is that the print spooler service stalls. The print spooler's job is to spool, or write the print jobs to the disk, and then once the print device is ready, it will send that job to the print device. If the spooler stalls, the jobs won't print, but you can't delete them from the queue.
When you have a situation like that-- the job won't print, can't delete it from the queue--that is a stalled print spooler. Go right into Services, restart the Print Spooler service, everything will start working again. Let's say that you find out that it's a particular driver that's constantly stalling your print spooler. You can go in and set up that driver to run in an isolated environment. That way, when the driver crashes, it's going to stall its environment, but it's not going to affect any of the other print drivers. I think that's really cool. Some people say to me, "Shad, why wouldn't I run every driver in isolation, if it's that great?" Every driver that you run in isolation is going to take up more resources on the print server. Practically speaking, unless you have a flaky print driver or you're troubleshooting, you don't want to necessarily configure every driver to be isolated.

Easy Print Driver
2:58-3:07
The Easy Print driver is used inside of Remote Desktop Services. Basically, what it does is allow somebody that's in a Remote Desktop Session to print to their local print device.

7.2.2 Using Print Management

Using Print Management
0:00-0:11
In this video, we're going to take a look at some of the features of print management. I'm going to go up under Tools and click on Print Management. First thing we see up here are Custom Filters, and that's exactly what they are.

Custom Filters
0:12-0:32
They let me filter according to criteria. For example, in the All Drivers filter, it's showing me all the drivers.
If I have printers and they enter the Not Ready state, they would drop into this filter. If they actually have jobs, then it would drop into this filter.

Creating a Custom Filter
0:33-0:38
I can create my own custom filter if I want to. It could be either Printer Filter or Driver Filter.

Filter Criteria
0:39-1:00
Once I get into this page, I can define my filter. I can say, OK, I'm interested in a situation where the Queue Status is exactly Toner/Ink Low.
I can also add further criteria if I want to, or I can just leave it at this.

Notification
1:01-1:45
When you set up filters, you also have the option to set up a notification. I can either have it send me an e-mail when a printer drops into this filter, or have it Run a script, or both. This is pretty nice. I can come in here, and if I do have somebody centrally for this print server who's in charge of toner, you can send an e-mail once any of those print devices get to low ink, and then they know to go ahead and swap out the ink or the toner cartridge.
If you want to send a notification on existing filter--say, you want to be notified anytime a printer is in the Not Ready status--you just right-click and go to Properties, Notification. Down here under my Print Server we're going to see Drivers, any drivers that we have installed.

Print Servers
1:46-2:03
We've got different Forms that can be used in the properties of the printer--Ports that are available to this print server, and then, Printers I've actually installed.
Yet, print drivers are famous for being flaky.

Print Drivers
2:04-2:43
I've seen it often in my career where I have two users--they have identical workstations; they have identical local print devices--but user A's machine wants to use the printer that comes down from the manufacturer's website and B's machine works fine with the printer that comes out of the operating system.
On computer B I just Next, Next, Next, I let it grab the print driver right off the hard drive, works fine. When I do that on A, the print device is not happy. I have to go download the driver. Print drivers can be very, very flaky. If you have a situation where you have drivers that are crashing or causing problems on the print server-- and that's a fairly common problem with Microsoft Print Servers--whether the print server is a client machine, or it's an actual server.

Drivers that are Crashing
2:44-3:10
Technically, the print server is whatever computer sends the job to that physical print device. These are pretty famous for stalling, and you know you have a stalled print spooler when you go into the printer, and inside the queue, you can't delete jobs, but they won't print.

Stalled Print Spooler
3:11-3:51
The jobs aren't printing, but I can't delete them-- that's almost always a stalled print spooler. What I would need to do is restart the Print Spooler Service. For that I would actually need to go into Services and I would just right-click Print Spooler, Restart. Much better than rebooting the server, faster and less risky. Now suppose you find out that, as a matter of fact, the print spooler keeps stalling, but the problem is this driver.

Running Unstable Print Drivers in Isolation

3:52-4:48

Maybe I have 40 drivers in this window, but only one of them is causing me a problem.

If you have unstable print drivers, you can run them in isolation, which means they'll run in their own virtual print server. When they stall, they'll stall the virtual print spooler, but it won't affect any other drivers. It will affect all the print devices that use this one driver, but just those print devices. Those print devices are still going to go down, but now you've got more time--or, less aggravated people as you go in and you restart your Spooler Service.

I would come in, I would right-click my driver, set Driver Isolation, and then set it to Isolated, and that will put it into its own little environment, and when it crashes, it just crashes those print devices. Those are some of the cool things about Print Management.

Summary

4:49-5:04

The nice thing too is, this is available on Windows 7 and Windows 8. If you have a client or personal computer and you want to use this to set up Notifications or Driver Isolation, you absolutely can.

7.2.3 Print Management Facts

Print Management, a utility for managing print servers, is available with Windows servers including Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012. Beginning with Window 7, Print Management is also available with client operating systems. Print Management automatically detects the printers on the same subnet as the print server. When printers are detected, Print Management automatically:

- Installs drivers
- Establishes the queues
- Shares the printers

The Print Management console allows you to centrally manage print server, print devices, and print drivers. Be aware of the following features of Print Management.

Feature	Description
Custom Filters	Custom Filters allow you to set conditions on which to filter. Default filters in Print Management are: • All Printers • All Drivers • Printers Not Ready • Printers with Jobs When you create a Custom Filter, specify if the filter is a printer filter or driver filter. Then set the filter criteria based on predefined fields, conditions, and values. You can also set up notifications to send an email notification or run a script when the criteria for the filter have been met. To create a notification on a default filter, use the **Notifications** tab in the Properties for the filter.
Printer permissions	Printer permissions control the actions that users can perform on a printer. Printer permissions apply to both local and shared printers. Beginning with Windows Server 2008 R2, administrative printer tasks, such as Manage Printer or Manage Documents, can be delegated to non-administrative users without the security risk involved in granting system administrative rights. In addition, printer permissions set on the print server become the default printer permissions for each new printer created on that print server.
Printer pooling	Printer *pooling* uses a single printer object to represent multiple print devices.

	• Printer pools speed printing by reducing the time that documents spend waiting for a free print device. • Printer pools simplify printer administration because you manage multiple devices through a single printer object. • With printer pooling, users send print jobs to a single printer object. • When creating a printer pool, all print devices in the pool must use the same printer driver. • Pooled print devices should be in the same physical location.
Multiple printer objects	You can use multiple printer objects for a single physical print device to customize access to the printer based on job roles. To configure multiple printers: 1. Create multiple printer objects, one per group or user with distinct access. 2. Configure permissions to restrict access for each printer. 3. Fine-tune access by editing the Advanced properties for the printer to modify priority (99 is the highest) and restricting printer availability times.
List in Active Directory	When you share a printer, you can choose to list the printer in Active Directory. Listing the printer in Active Directory makes the printer name and its characteristics appear in Active Directory. Users can then search Active Directory to find the printer by name or by special features (such as location or color support). Listing the printer in Active Directory does not automatically add a printer to client computers. Users must still connect to the printer.
Deploy with Group Policy	When you deploy a printer using Group Policy, printer objects are automatically created on client computers that point to the deployed printers. • You can deploy printers using the Print Management snap-in or by editing the Group Policy Object (GPO) in the Group Policy Management snap-in. • Deploying the printer to a computer adds the printer to the computer, regardless of what user is logged on. Deploying the printer to users adds the printer for the user, regardless of which computer they log on to. • Machines that run Windows 2000 support only per user connections. Machines that run XP or later support both per user and per machine connections. • For pre-Vista computers, you must run the **PushPrinterConnections.exe** utility to have the workstation pull printer information from the Group Policy. Run this on a periodic basis so that changes to the deployed printers are updated on the

	client. As a best practice, run **PushPrinterConnections.exe** as a script in the same GPO where you have deployed the printer.
Export/import printers	Using Print Management, you can export and import printer settings including print queues, port settings, printer settings, and language monitors. Once you export the settings to a file, you can import them to another print server. The import options are as follows: • **Import Mode** determines the action to take if a print queue exists at the destination computer. • **List in the Directory** determines if the import print queue is published in Active Directory directory service. • **Convert LPR Ports to Standard Port Monitors** determines if Line Printer Remote (LPR) printer ports from the printer settings file are converted to the faster Standard Port Monitor standard. You can also perform printer migration from the command line using the **Printbrm.exe** command.
Manage print drivers	When you configure a shared printer, computers that connect to the shared printer automatically download the necessary print driver. When managing print drivers on shared printers, keep in mind the following: • On the print server, install additional drivers that correspond to the hardware platform for the installed printers. You can add drivers by editing the printer object, or through the print server properties. • Windows NT and later systems will download any new or updated drivers associated with the shared printer. When you change the driver on the shared printer, the clients will download the new driver. • Windows 9x systems (95/98/ME) will not download drivers automatically. You will need to manually configure and update drivers.
Manage documents in the print queue	For documents in the print queue: • Users can delete or pause their own print jobs. Those in the Print Operator group can manage all documents in the queue. • Pausing the queue prevents any document from printing. • To change the order that documents print, change the document priority. Documents with a higher priority number print first. • Use the Reliability and Performance Monitor to gather statistics about print queues. Statistics you can monitor include: Jobs, jobs spooling, and job errors Not ready and out of paper errors Total bytes, jobs, and pages printed

	• Use printer filters in Print Management to view the current state of printers and print queues and to configure e-mail notifications for printer error conditions. You can also use alerts in Reliability and Performance Monitor to receive notifications about print queue statistics. Print Management allows you to track more error states than Performance Monitor, but Performance Monitor lets you view performance statistics that are not available through Print Management. Performance Monitor also lets you capture information about print queues and save that data to a file. Print Management lets you see the current status or send e-mail notifications only.
Location-aware printing	Location-aware printing sets a default printer for each network connection on a mobile client. Location-aware printing: • Automatically adjusts according to the location and selects the correct default printer for that network. • Works only on portable devices (devices with a battery), not desktop computers. • Is managed through the **Manage Default Printers** dialog in Devices and Printers. You select the network and identify the default printer for that network. To configure a default printer for a wireless network, you must have connected at least once to the wireless network. • Is disabled by choosing the **Always use the same printer as my default printer** option. • Does not work through Remote Desktop connections because the mobile device must be physically connected to the network.
Print driver isolation	Driver isolation allows printer driver components to run in an isolated process separate from the printer spooler process. Isolating the printer driver: • Increases print server reliability. • Allows you to test and debug new drivers by isolating them. • Can be used to identify printer drivers which are causing spooler failures. Running print drivers in isolation requires more resources and should be used only for testing, debugging, and identifying problem print drivers.

Client-Side Rendering (CSR)	Client-Side Rendering (CSR) renders or spools print jobs on the client machine to reduce print processing times. CSR: • Reduces the CPU and memory load on print servers. • Reduces network traffic. • Is enabled by default. It can be disabled through the **Sharing** tab in the printer properties.
Easy Print driver	The Easy Print driver is used in Remote Desktop Services. Easy Print driver allows a user in a remote desktop session to print to their local print device without having to install a print driver.

When you configure printing, you create a logical printer object that references a print device or points to another logical printer on the network. The following table lists the configuration choices for each type of printer.

Print Device Location	Printer Type	Port Type
Connected to the LPT, USB, or COM port of the local computer	Local	LPT, USB, or COM
Connected directly to the network through a NIC connected to the printer	Local	TCP/IP (identify the IP address of the print device NIC)
Connected to a shared printer that is configured on a remote computer	Network	UNC path (*computername**sharename*)

8.1 Group Policy Foundation

As you study this section, answer the following questions:

- How does inheritance affect Group Policy settings?
- If a setting is configured in a GPO linked to the domain and a GPO linked to an OU, which setting will be in effect?
- What is the difference between using a starter GPO and copying an existing GPO?

After finishing this section, you should be able to complete the following tasks:

- Create and link GPOs to appropriate objects.
- Create a starter GPO.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  - Manage Group Policy Objects (GPOs)
    - Create and Link a GPO
    - Create a Starter GPO
    - Modify GPO Links

This section covers the following 70-410 exam objective:

- 601 Create Group Policy objects (GPOs).
  This objective may include but is not limited to:
  - Manage starter GPOs
  - Configure GPO links
  - Configure multiple local group policies

8.1.1 Local Policies

Local Policies
0:00-0:15
Let's talk about Local Group Polices. Group Policy is a way to make changes to a work station without having to visit it, and we're really looking at centralized administration. I want all my computers to look identical, and if possible, I don't want to have to visit them.

Local Group Policy
0:16-0:42
A Local Group Policy affects only that particular machine. You might be thinking, why would we want a Local Group Policy? I tend to think of it like this: the settings that you often configure on a work station, those are in Control Panel, and Control Panel is obviously named; it's the place where we go-- it's the Control Panel of the computer. These are settings that it's okay for a regular end user on their home computer to go in and make changes, and it's probably not going to hurt that machine.

Group Policy
0:43-0:55
Things that get a little bit more complex and really should only be addressed by an administrator are done with Group Policy. The idea being that you've got to know a little something to be able to get into Group Policy to make the change.

Registry Edits
0:56-1:12
Settings that are really obscure-- meaning, Microsoft doesn't really want you changing them unless you absolutely have to--those are generally done with Registry Edits.
My philosophy on the Registry is I'm not going in there to touch anything unless I know for a fact that this is how I'm going to solve a specific problem.

Local Policies for Kiosk Machines
1:13-2:21
Local polices are generally not used that much unless you have a standalone machine--some kind of machine in your environment that's in a work group and not a member of the domain. Usually, we call these Kiosk machines. Let's say you have a computer in the lobby of your company. You really don't want anybody coming in and seeing any information about the domain. You leave it as a standalone computer, but you do want to configure it using Group Policy. In that case you would use a Local Group Policy.
The only place we would use this in a domain is if we have something that really just applies to that machine. There's no way that we can tailor the policy from inside Active Directory to just hit that machine. That, to me, would be a last resort. I'm always looking to do things centralized from Group Policy if it's a domain machine, not only because it makes it less work on me, but also because it's better documentation. When I win the lottery and I don't come back to work tomorrow, the next person coming in will be able to see exactly what I've done just by looking at the Group Policy.
Sometimes you can go through and set up the Local Policy and build that right into your image, so that all of the computers have the same Local Policy.

Multiple Local Group Policies
2:22-3:13
The only other thing we need to know about Local Group Policies is, starting with Windows Server 2008 R2 and the Windows 7 client, we can actually have multiple Local Group Polices. Traditionally, before that, there was just one. If I put in settings that would, let's say, restrict the user, and say I want to take away the Start button, I want to disable Control Panel, I want to go ahead and make sure that they can't create any additional icons on the desktop. Well, that would be in effect for every user including the administrator, which now makes my job of supporting that computer much more difficult.
With multiple Local Group Polices, I can have one policy that affects the computer, meaning anybody logs in to that computer and set up the computer half of Group Policy. But, I can have multiple Local Group Polices that affect different users. I can make a Local Group Policy that affects only administrators, only non-administrators, or specific users and groups.

Configuring Local Policies
0:00-0:18
In this video, we're going to take a look at configuring Local Policies. To get into the Local Policy on any computer, you want to go ahead and type in "gpedit.msc". I'm going to right click "gpedit" and choose Run as administrator. This is showing me the Local Policy, and it's the Local Computer Policy.

Local Computer Policy
0:19-0:25
It affects just this computer.

Computer and User Configuration Settings
0:26-0:48
The Computer Configuration settings affect the entire computer. The User Configuration settings right now affect all the users. I can go through and do all the things that I would normally do with group policy in here. I've got Windows settings. I've got some security settings. I've got administrator templates, so on and so forth.

Multiple Local Policies
0:49-1:15
The quickest way to get into the Local Policy--new starting with Windows Server 2008 R2 Windows 7, which means it's effective for Windows Server 2012 and Windows 8--we have the ability to have Multiple Local Policies. To work with those, you have to go in in a different way. You can't use gpedit.msc. So I'm going to go ahead and open up Microsoft Management Console, MMC, and I'll Run that as an administrator.

Microsoft Management Console (MMC)
1:16-2:38
MMC was intended to be a single administrative tool that I could use to manage anything on the computer. So when you first open it up, it doesn't have any snap-ins. I've got to go up under File and Add/Remove Snap-in to add tools to make it functional.
For Local Group Policies, I want to add the Group Policy Object. It comes up, and by default it's going to add that Local Computer Policy that we were just in. I'll go ahead and add that. The Multiple Local Policies are added a little bit differently. If I add it again and I hit Browse, I can select Group Policies that affect individual users or groups of users. Right now this computer only has one user on it, the Administrator. I could do a policy that just affects all administrators. I could also do a policy that just affects non-administrators.
You'll see the Local Computer Policy is exactly what we saw with gpedit.msc. Because these two policies just affect users, they only have the user configuration. Here's why this is so cool.

Benefits of Multiple Local Policies
2:39-3:48
When we only had one computer policy, let's say I have a kiosk computer and I decide that this kiosk computer, it's running in the cafeteria; all I really want it to do is be used to browse the internet for when people are at lunch. I go into Administrative Templates and I start to set up the Control Panel so that there is no Control Panel. Prohibit access to Control Panel. Maybe I go into the desktop and I hide and disable all the items on the desktop, Remove the Computer icon. I get rid of Documents. I turn off everything except Internet Explorer. Well, with just one local computer policy, those particular settings are going to affect everybody who logs in. Now, as an administrator, I've got to deal with this computer that's completely locked down. Now with Multiple policies I can go through and turn all of these things off for Non-Administrators and yet leave them on for Administrators.

Summary
3:49-4:34
Some of the specific sections of Group Policies are covered in other videos, but our goal with this video is just to talk about the Local Policy, how we get into it, and how we would use Multiple Local Group Policies. To get into your Local Computer Policy, that's your gpedit.msc. If you want to use Multiple Local Group Policies, I want to go into MMC, Add or Remove the snap-in, and then browse when I'm in the Group Policy Object. Here's another cool thing about MMC before we go. If this is what I normally work in, I can save the console, and then I can reopen it in the future. That's how we configure Local Policies.

8.1.3 Group Policy Processing

Group Policy Processing
0:00-0:02
We're going to talk about Group Policy processing order.
Group Policy
0:03-0:22
Group Policies are a way to make configuration changes to the workstation without having to visit it. What we want to do is go into Active Directory and set up a Group Policy that will make some kind of configuration change, so that the workstation is more secure, or runs better, or whatever it is we're trying to do.
Security
0:23-0:54
Generally speaking, when you're setting up security in a network, the rule of thumb is it's your job as the network administrator to take away as many rights as possible from the end user. It's their job as the end user to cry and get those rights back, so they should end up with exactly what they need to do their work--no more, no less. You don't want to go ahead and go to town with Group Policy, because I've actually seen setups where there are so many restrictions you can't even use the computer.
LSDOU
0:55-1:56
In order to know what the end result will be for the user, you need to know the order in which the policies are going to run and how they're going to be processed. There's a little acronym for that called LSDOU. Let's take a look at a diagram I've put up on the whiteboard for you. First of all, any particular policy usually has three settings: enabled, disabled, or not defined. When I think of Group Policies, I think of light switches.
Basically, if you think about a light switch, the last person that walks through the door and flicks the switch is the one that wins. If you walk in last, you turn it off, the light is off. You walk in last, you turn it on, the light is on--same with Group Policy. If the setting is enabled, it turns the setting on. If it's Disabled, it turns the setting off. Not defined means don't make a change, so it will stay whatever is the default, or whatever has been set by a Group Policy that ran before that Group Policy.
When you're working with Group Policy, also make sure you notice whether it's a positive policy or a negative policy.
Positive/Negative Policy
1:57-2:18
Sometimes, there's double negatives. For example, there might be a policy that says, "Disable the network adapter," so that when you disable the Disable the network adapter, what you really have is a network adapter that's turned on.
Levels of Group Policy
2:19-2:22
There are four levels of Group Policy.
Local Policy
2:23-2:48
The very first Group Policy that runs is the Local Policy. Because this runs first, it's the least powerful. The idea is that I am going to put specific settings in there that really apply to that particular computer only, and presumably, it's not going to be contradicted later on by any of these. Again, we generally don't work with that very much. We would prefer to work in Active Directory.
Site Policy
2:49-3:16
The next policies that we'll run are any policies attached to the Site, and we have a different lesson where we talk about exactly what Sites are. They're the geographic location, usually all in one building. Anything that needs to happen while they're in that particular location, maybe there's a specific wireless network, that will be done in the Site Policy, or maybe they need a setting. They only need that setting when their laptop is at work. I could attach it to the Site.
Domain Policy
3:17-3:26
The next policy to run is the Domain Policy. We put settings in the Domain Policy when we want them to affect everybody in our domain.

Organizational Unit
3:27-4:43
The last policies to run are policies that are linked to Organizational Units. If there's nested Organizational Units, like here I have an OU called Sales, here's a nested OU inside of Sales called Desktops, first the policies at Sales will run, then the policies at Desktops will run.
The idea is, Site we're really only going to use it if it's a geographic thing. That's not that common. We've got our base polices that set up our environment for the whole domain, and then as we get closer and closer to the user or the computer objects, we're setting up policies that are going to build the ultimate environment that we want to have happen.
With that having been said, you want to try to have the policies build on each other. Even though we talk about one policy overriding another policy, in an ideal situation, we don't need to do that. We set up our generic policies at the domain, maybe nobody can use removable hard drives, you want certain things on the desktop, and then our OUs are done with the logical structure of our organization. There may be salespeople who need specific things, sales desktop needs specific things, and they build together.

Policies that Reverse Each Other
4:44-5:12
It is possible that you will have policies that reverse each other. Let's say we do have a policy at the domain that says no removable hard drives of any kind, but we want our LAN administrators to be able to use flash drives to run diagnostic tools. Then we could go ahead and make a policy that would allow removable drives, but would run later on in the Group Policy processing order, and only affect our network administrators.

Link Order
5:13-5:58
I have multiple Group Policies linked to an OU, so here I have created a Sales OU and I've got Group Policy 1, 2, and 3. When we go into the demo, we're going to see on the tab of the OU exactly where it'll be listed out. These policies are a process from the bottom up, so first GPO3 would run, then 2 would run, then 1 would run. The policy that's linked highest is the one that's going to win, so to speak.
If you come into a scenario where you have an issue where one policy is overriding another policy and they're both linked to the OU, in that case, you're going to change the link order. That's what this is called.

Group Policy
5:59-6:06
When we get into the demo, you're going to see that Group Policy really has two halves: a computer node and a user node.

Computer Policies Order
6:07-6:40
The computer policies run when the computer boots, so our processing order is going to be followed when that computer boots.
The client1 in the Sales Desktops OU boots up, all of the computer half of the Local Policy will run; the computer half of the Site Policy will run, the computer half of the Domain Policy will run, the computer half of any policies at Sales will run, and the computer half of any policies at Desktops will run, and then I see "hit Ctrl + Alt + Delete" to log in.

User Policies Order
6:41-7:04
The user half of the policies run when the user logs in. As I'm logging in, the user side of any Local Policies, user side of any Site Policies, user side of any domain policies, and then the user side of any OUs in which my user account is located. That's how we build the ultimate overall environment for the user.

Summary
7:05-7:23
What you really need to have a grasp of is the order in which these will be processed, which is our LSDOU. First the Local runs, then the Site runs, then the Domain, and then the OUs.
If you have a good grasp of the processing order, you should be able to troubleshoot problems fairly effectively.

8.1.4 Linking Group Policy Objects

Linking Group Policy Objects
0:00-0:24
In this video, we're going to take a look at linking Group Policy Objects. I need to go into my Group Policy Management console. I'm going to click on Tools > Group Policy Management. Let's make this a little bit bigger, we can really see what's going on. The Local Policy is kept at the local workstation. From an Active Directory standpoint, I have three levels at which I can link policies.

Site Policy
0:25-1:02
The order in which they're processed is this: any policy that's linked to a Site, and by default, I can't see the Sites, so I'd have to click on Show Sites, put a check mark in the Sites that I want to see, and then I can go ahead and manage Group Policy. I would attach a Group Policy to a Site if I want that policy to be in effect if the computer or the user is at that physical location, or maybe I only want it to be in effect when they're connected to the work network. Maybe I'm going to roll out wireless settings, but I don't want to override their wireless settings when they're at home, so I'll link it to my Sites. It's only going to take effect when they're at a company site.

Domain Policy
1:03-1:12
The second level that's processed is the domain. Anything linked to the domain pretty much is going to affect all my users unless I do some funky things and mess around with it.

OU Policy
1:13-1:36
I'm linking policies to the domain when I want it to affect just about everybody, and then I can also link it to the OU if I want it to just affect these users or computers inside of that OU. If I have an OU inside of an OU, anything attached at the parent OU will be inherited by the child OU. If I attach a policy at sales, it's going to be in effect for managers as well.

Active Directory Policies
1:37-1:42
Let's take a look at the policies that come with Active Directory. Right off the bat, you can see I have a Default Domain Policy and this is linked to my domain.

Default Domain Policy
1:43-2:14
Basically, if I need to go through and make changes, I can use this policy. They don't necessarily recommend that you mess around with this policy. You certainly can, but it's probably better to make your own. On the other hand, the more policies you have, the more policies the workstations have to process. It might slightly down their boot. If you don't need a separate policy, make as few as you actually need. If you can combine multiple settings into one policy, that's the best way to go.

Password Policy
2:15-2:25
With a Default Domain Policy, the most important piece is the Password Policy, because this is the only Password Policy that's going to take effect unless I create Password Settings Objects.

GPO's
2:26-2:55
Notice I can right click my domain and I can link GPOs here. There's no such thing as a GPO up at the forest level. If I have a GPO that should be in use in multiple domains, I'm going to create it in one domain, export it, give it to the network administrators in the other domain, let them import it and use it in their domain. It's not advised to link GPOs from another domain and pull them over. I also am not able to attach GPO at the forest level. You want to make sure that you know that.

Default Domain Controllers Policy
2:56-3:22
In addition to the Default Domain Policy, I have a Default Domain Controllers Policy which is linked to the Domain Controllers organizational unit by default. Basically, what that does is tighten up security on the Domain Controllers. For example, it says only administrators can sit down at the computer and log in, have an interactive log on. Again, if I need to make any changes, I certainly can do that, but for the most part, I leave these alone. If I need to reset them, if I completely messed them up, you can run the dcgpofix command, and that will reset them to exactly what they look like when you installed Active Directory.

dcgpofix Command
3:23-3:33
Example Linking
3:34-3:50
Let's go ahead and create some policies and link them. I'm going to right click my domain and I'm going to create a GPO in this domain and link it here. We're going to go ahead and Block USB Drives. There are Starter GPOs that exist there.
Starter GPO
3:51-4:56
In the Starter GPO set, you can create them. They're supposed to be a baseline template for security. What you could do is--if you want to have a security standard for your organization-- you can set up the starter GPOs and then tell the LAN administrators when you make a new GPO you want to base it off of the starter GPO. They're not really heavily used, but they're certainly out there. I'm not going to copy any GPO, I'm just going to say ok. Now you can see this is just a link. It's that little arrow there. The GPO itself actually lives inside this Group Policy Objects container. If I delete this and it tells me this is just a link, but if I edit it from here, even though I'm editing it from the link, it's going to change the actual GPO which affects everywhere it's linked. That's essentially what this message says. If I delete the link, you can see that the link goes away but the GPO still continues to live in Group Policy Objects. I would need to come and delete it out of here if I want to completely get rid of it.
Link Existing GPO
4:57-5:36
If I want to link an existing GPO, just click "Link Existing GPO", and we'll relink lock USB drives.
Let's take a look at what happens if I create another GPO. We're going to just say for the sake of it that the users in the Sales OU and in the Managers OU are allowed to have USB drives. I'm going to create another GPO that says Allow USB drives. I realize I haven't gone in and actually blocked them. For our discussion, we're not looking at any of the contents of the policy. We're just looking at linking the policies.
Group Policy Inheritance
5:37-6:21
If I ever have any questions about what's going to happen, I can go over to Group Policy Inheritance and it tells me what's going to take effect. This list is processed from the bottom up, so whichever one has the lowest number has the highest priority. First, it's going to process the block USB drives policy. Then, it's going to process the Default Domain Policy. Last, it's going to process the Allow USB Drives policy, which means that effectively, the users inside of Sales and inside of Managers are going to be allowed to use their USB drives. I can check that just by coming down to Managers, and you can see that it's the same if I have multiple policies on the same OU, so let's add a policy, so I now have a policy called Redirect Documents.
Multiple Policies
6:22-7:05
If I click over on Linked Group Policy Objects, in an ideal world, these policies should not conflict with each other. Make no mistake, very rarely do we have a situation where we're going to block something up at the domain level, but then unblock down here. You should build domain policies that affect everybody, and then as you add policies to OUs, it's going to add into them and set up your environment the way it needs to be set up. It's not trying to turn stuff off, on, off, on, like that. If you do have a situation where for whatever reason there are settings in these two that conflict, let's say you have a situation where you say, "I know if Allow USB Drives was processed first, and then Redirect Documents was processed, everything would work great, but as it stands, the computer's going to process Redirect Documents and then Process USB Drives." For whatever reason, it's creating a problem.
Adjust Link Order
7:06-7:56
In that case, you can adjust the link order because they're both linked at this level. I can just click the up arrow here and rearrange them so that they come in the proper order. Again, if you're doing this all day long, you need to redesign Group Policy. They're meant to work together, not to override each other.
Summary
7:57-8:08
That's how we go through and we manage our Group Policy links, creating and linking them where we need to. Remember, they live in the Group Policy Objects container, so if I delete a link, it's not going to do anything to the policy.

8.1.5 Group Policy Categories

GPOs contain hundreds of configuration settings. The following table describes common settings.

Setting Category	Description
Account Policies	Use Account Policies to control the following: • Password settings • Account lockout settings • Kerberos settings Account policies are in effect only when configured in a GPO linked to a domain.
Local Policies/Audit Policy	Use Audit Policy settings to configure auditing for events, such as log on, account management, or privilege use.
Local Policies/User Rights Assignment	Computer policies include a special category of policies called *user rights*. User rights identify system maintenance tasks and the users or groups who can perform these actions. Examples of user rights include: • Access this computer from the network (the ability to access resources on the computer through a network connection) • Load and unload device drivers • Back up files and directories (does not include restoring files and directories) • Shut down the system • Remove a computer from a docking station
Local Policies/Security Options	Security options allow you to apply or disable rights for all users to whom the Group Policy applies. Examples of Security Options policies include: • Computer shut down when Security event log reaches capacity • Unsigned driver installation • Ctrl+Alt+Del required for log on
Registry	You can use registry policies to: • Configure specific registry keys and values. • Specify if a user can view and/or change a registry value, view sub-keys, or modify key permissions.

File System	Use File System policies to configure file and folder permissions that apply to multiple computers. For example, you can limit access to specific files that appear on all client computers.
Software Restriction Policies	Use software restrictions policies to define the software permitted to run on any computer in the domain. These policies can be applied to specific users or all users. You can use software restrictions to: • Identify allowed or blocked software. • Allow users to run only specified files on multi-user computers. • Determine who can add trusted publishers. • Apply restrictions to specific users or all users.
Administrative Templates	Administrative templates are registry-based settings that can be configured within a GPO to control the computer and the overall user experience such as: • Use of Windows features such as BitLocker, Offline files, and Parental Controls • Customize the Start menu, taskbar, or desktop environment • Control notifications • Restrict access to Control Panel features • Configure Internet Explorer features and options
Starter Group Policy Objects	Starter Group Policy Objects, referred to as Starter GPOs, allow you to store a collection of Administrative Template policy settings in a single object. • When you create a new GPO from a Starter GPO, the new GPO has all of the Administrative Template policy settings and values that were defined in the Starter GPO. • You can easily distribute Starter GPOs by exporting and then importing them to another environment.

8.1.6 Group Policy Facts

A *policy* is a set of configuration settings applied to objects such as users or computers. Group policies allow the administrator to apply multiple settings to multiple objects within the Active Directory domain at one time. Collections of policy settings are stored in a Group Policy object (GPO). The GPO includes registry settings, scripts, templates, and software-specific configuration values.

The following table identifies tasks for managing GPOs.

Task	Description
Creating local GPOs	A local GPO is stored on a local machine. Computers that are not part of a domain use the Local Security Policy settings to control security settings and other restrictions on the computer. To edit the local Group Policy, enter **gpedit** at the command line. Beginning with Windows Server 2008 R2 and Windows 7, a local computer can have multiple local Group Policies: • One Group Policy that affects the computer • One or more Group Policies that affect users To create or edit multiple local Group Policies, you use the Microsoft Management Console (mmc): • Enter **mmc** at the command line to launch the Microsoft Management Console. • Add the **Group Policy Object Editor** snap-in from the File menu. By default it will add the Local Computer Group Policy. • Select **Users** to edit Local Group Policy for specific users on the computer. You can save the Group Policy Object Editor console to allow for easy access in the future.
Assigning GPO permissions	Group Policy permissions control the operations that users can perform on the GPO as well as the application of the GPO to the user. • To apply settings to a user, the user must have the **Allow Read** and **Apply Group Policy** permissions. • By default, each GPO grants the Authenticated Users group (basically all network users) the **Allow Read** and **Apply Group Policy** permissions. This means that, by default, GPO settings apply to all users. • Permissions also control who can edit Group Policy settings and manage the GPO.

Linking GPOs	GPOs can be linked to Active Directory sites, domains, and organizational units (OUs). Use the Group Policy Management console to link Group Policy. • A GPO applied to an OU affects the objects in the OU and sub-OUs. • A GPO applied to a domain affects all objects in all OUs in the domain. Built-in containers (such as the Computers container) and folders cannot have GPOs linked to them. When linking Group Policies • The Default Domain Controllers Policy is linked to the domain controllers OU by default. This policy increases security of the domain controllers. You can run the **dcgpofix** command to restore the original settings of the Default Domain Controllers Group Policy. • On the **Linked Group Policy Objects** tab you can change the link order of Group Policies. • The **Group Policy Inheritance** tab lists the order in which Group Policies will be applied. The policies are listed in reverse order of precedence, meaning that the last policy on the list--the one with the highest precedence number--will be applied first. • To delete a Group Policy, you must delete it from the Group Policy Objects container.
Using Administrative Templates	You can use Administrative Templates to create Group Policies to manage Microsoft Office or in-house applications. File types for Administrative Templates use an XML-based file format that allows multi-language support and version control: • .admx files are the Administrative Template files and require Windows Vista or later to edit. • .adml files contain the language-specific Administrative Template files. .adm files are the pre-XML format used for Administrative Templates. This older format is still usable in Windows Server 2012.
Using a central store	When you use Administrative Templates, the policy is stored locally and the settings are saved to Group Policy on the domain controller. The central store allows Administrative Templates to be available to be edited by other domain administrators. • Group Policies are kept in SYSVOL, a share that is created when you install Active Directory. All domain controllers in the domain have a replicated copy of SYSVOL. • To create a central store:

	Create a folder named PolicyDefinitions in file:\\FQDN\SYSVOL\FQDN\. For example: **\\Northsim.com\SYSVOL\Northsim.com\PolicyDefinitions** Copy the contents of the local PolicyDefinitions folder to the PolicyDefinitions folder on SYSVOL. The path of the local PolicyDefinitions folder is typically: **C:/Windows/PolicyDefinitions**

Keep in mind the following about GPOs:

- If possible, combine multiple settings into one Group Policy. Reducing the number of Group Policies that require processing reduces boot and logon time.
- The Default Domain policy contains the only password policy that is going to take effect, unless you create a password settings object (PSO).
- GPOs do not exist at the forest level. To enforce a GPO in multiple domains, create the GPO in one domain, and export it and then import it into other domains.

Each GPO has a common structure, with hundreds of configuration settings that can be enabled and configured. Settings in a GPO are divided into two categories:

GPO Category	Description
Computer Configuration	Computer policies (also called *machine policies*) are enforced for the entire computer and are applied when the computer boots. Computer policies are in effect regardless of the user logging into the computer. Computer policies include: • Software that should be installed on a specific computer. • Scripts that should run at startup or shutdown. • Password restrictions that must be met for all user accounts. • Network communication security settings. • Registry settings that apply to the computer (the HKEY_LOCAL_MACHINE subtree). Computer policies are initially applied as the computer boots, and are enforced before any user logs on.
User Configuration	*User policies* are enforced for specific users. User policy settings include: • Software that should be installed for a specific user. • Scripts that should run at logon or logoff. • Internet Explorer user settings (such as favorites and security settings).

	• Registry settings that apply to the current user (the HKEY_CURRENT_USER subtree). User policies are initially applied as the user logs on. They often customize Windows based on user preferences.

All computer policies run before the user policies run.

8.2 Group Policy Management

As you study this section, answer the following questions:

- What are the advantages of the .admx file format?
- What is the Administrative Template central store? What advantages do you gain by enabling the central store?
- How is the **Block Inheritance** setting affected by the **No Override** setting?
- How does *loopback processing* affect computer settings?

After finishing this section, you should be able to complete the following tasks:

- Centrally manage administrative templates using the central store.
- Configure the scope of Group Policy objects.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Manage Group Policy Objects (GPOs)
  - Modify GPO Links

This section covers the following 70-410 exam objective:

- 601 Create Group Policy objects (GPOs).
  This objective may include but is not limited to:
  Configure a Central Store
  Configure security filtering

8.2.1 Central Stores

Central Stores
0:00-0:10
Let's talk about creating a central store. Group Policy is a way to make changes to the workstations from a central environment. There are parts in Group Policy that come when you install Windows or Active Directory.

Customizable Group Policy
0:11-0:52
It was also made to be extensible, so the idea was that you could create Group Policies of your own to manage in-house applications or settings. There's a bunch of Group Policies that you can add in for Microsoft Office, pretty much any version, anything you can imagine in Office can be configured using Group Policy, or you could have your programmers create Group Policies for in-house applications if there's settings that need to be applied to the workstation.
These customizable Group Policies are all kept in the Administrative Templates Node in Group Policy. If you add in new Administrative Templates, they get added in on the local workstation.

ADM File
0:53-1:00
Originally, these were called .ADM files. That would be the old style Windows Server 2003.

ADMX File
1:01-1:39
You may or may not remember when in Office, suddenly everything had an 'X' at the end. Before it was .doc, now it's DOCX. Here it's ADM, now it's ADMX.
The X, by the way, stands for XML. Everything went to XML so that everything would be compatible with browsers and would have a standard way of displaying documents. ADMX files have to be edited on Windows Vista or better, so Vista, Windows 7, Windows 8, great. We can't use Windows XP to edit ADMX files.
We still have compatibility with the old ADM file, so if your company has ADM files or custom applications, you can use those, but everything new will be ADMX.

ADML File
1:40-2:07
If you have language-specific administrative templates, those will be kept in ADML files. If you need French or Spanish, you need an ADML file. Those are the ones that are language specific, and inside the central store, you're going to create a folder with the correct code for the language-- like US English is en-us--but each language has a different correct acronym, and then you'll store your ADML files in the appropriate folder in the central store.

AMDX Files and the Central Store
2:08-3:06
By default, when I'm editing those AMDX files on my workstation, they're kept in a folder in the C: drive, or wherever I have installed Windows, usually at the C: drive. Inside of the Windows folder, there'll be a folder called Policy Definitions. If I add in a bunch of Administrative Templates for Microsoft Office, to start with, they're going to be kept on my local workstation. Now those settings will be saved up to Group Policy.
The problem comes in if another administrator needs to edit my custom ADMX template. When they go to open it up on their machine, they're not going to see that template. That's why we create a central store. Central store gives us a place to store these templates centrally so that no matter who opens up the Group Policy on whatever workstation, they're going to see all the administrative templates that I'm using in my organization.
The moral of the story is, if you're using custom ADMX templates, you need to create a central store.

SYSVOL
3:07-5:03
Group Policies are kept in the SYSVOL, which is a share that's created when you install Active Directory. All of the domain controllers in the domain have an exact copy of the SYSVOL, so when I connect up to a domain controller, I edit a Group Policy, the changes are stored in the SYSVOL and then if I were using Server 2012, the Distributive File System, DFS, takes care of replicating the contents of that share to all the other domain controllers. It's separate from Group Policy replication. The file replication, because these are files in a share, but that share is kept synchronized between all the domain controllers.
I want to create my central store up in the SYSVOL because number one, that's where Group Policy lives; and number two, that's a share that's already replicated to all the domain controllers, which means no matter which domain controller an administrator connects to, they're going to have access to those templates. We can access the SYSVOL by typing in \\ then the Fully Qualified Domain Name of our domain, \SYSVOL, then there's going to be the Fully Qualified Domain Name of the domain again.

Inside that folder, we need to create a PolicyDefinitions folder. That's really what the central store is. Once I've created that PolicyDefinitions folder inside of here, I go to my local PolicyDefinitions, and I'm going to copy anything that's in there up to the central store. Just so you know what the central store might look like in a real domain, supposing that my domain is named NorthSim, it would be \\NorthSim.com\SYSVOL\NorthSim.com\. Again, I ran out of space, so I'm going to roll it to the next line, PolicyDefinitions. So NorthSim.com, that's my Fully Qualified Domain Name, SYSVOL, NorthSim.com again, another slash, and then here's the Policy Definitions folder I'm going to create, copying everything in that local policy definitions up to the central store, and then I should be good to go.

Summary

5:04-5:45

My central store allows me to centrally store custom ADMX files, which are used to expand Group Policy by adding Custom Administrative Templates, either from Microsoft applications or for in-house applications, to make sure that those templates are available all throughout the domain, so that any administrator that opens up Group Policy to edit it will see all those new options.

You need to memorize the path for the Local Policy Definitions folder and the path where you're going to put it up on the SYSVOL. If you're not a great memorizer, at least know it's a folder called PolicyDefinitions, and I'm copying it from the local workstation up to some place in the SYSVOL.

8.2.2 Creating a Central Stores

Creating a Central Store
0:00-0:03
In this video, we're going to see how to set up a central store.
Administrator Templates (ADMX Files)
0:04-1:15
In Group Policy, there's a section of Group Policy that's extensible called the Admin Templates or Administrator Templates. I've made this very wide, but you can see that Administrative Templates are kept in ADMX files and they're retrieved from the local computer.
The idea behind it was this: everything in Software Settings is for rolling out software to computers. Windows Settings are settings for Windows, which are controlled by Microsoft. But Administrative Templates were intended to be settings that could be extended. For example, there are Administrative Templates that you can get from Microsoft Office that would add a whole bunch of settings into this node in Group Policy to set up Word, Excel, PowerPoint.
I can even go through if I have custom applications in my environment and have the programmer set up Group Policy templates in an ADMX file. The X is for .xml--they're .xml files--and that programmer can create an .xml file that I could then import into Group Policy and go through and control the settings for those custom applications.
Languages (ADML Files)
1:16-1:24
If I want to see the settings in Administrative Templates in different languages, we use ADML files. Just think "L for languages." Right now, my Administrative Template files--if I had any custom ones--are being retrieved from the local computer.
Location of ADMS Files
1:25-1:49
They're kept in the C: Windows PolicyDefinitions Folder. This is where the computer is currently pulling these ADMX files from. You see there's a bunch of them in there.
Location of ADML Files
1:50-2:16
Language-specific files--those ADML files--are kept in folders that are named for that specific language. This is English U.S. If you needed French, you could make a French folder, Spanish folder, whatever language you need you find out what the code is, and you create that folder.
If I add in a bunch of new templates that would be kept on the local computer, and I could go through and I could configure Group Policy.
The Need for Central Store
2:17-2:50
The problem comes in if one of my colleagues opens up Group Policy Management Editor on a different computer. Then he or she is not going to have the files that are on my computer.
When I add in extra ADMX files, it's not going to be sufficient to have them stored on the local computer. In that case, I need to create a central store. Central store is just a folder where the computer can go to pull these centralized ADMX files, so that no matter who opens Group Policy, we're all going to see the same templates.
Opening the SYSVOL
2:51-3:04
What I'm going to need to do is this: I'm going to need to open up the SYSVOL. The SYSVOL is shared; it's a share that's on all my domain controllers, and it's replicated between all the domain controllers.
Group Policies
3:05-3:17
To get into it, I go to // the name of my domain, SYSVOL, the name of my domain again, and this is where the actual Group Policies are kept. It's also where I would put logon scripts.
Default Domain Policy and Default Domain Controllers Policy
3:18-3:53
So, there are a couple of group policies that come with Windows. The Default Domain Policy and the Default Domain Controllers policy. If I look in Policies, there's two GUIDs in here. One of them is the Default Domain Policy; the other is the Default Domain Controllers.
As a bit of trivia, these numbers are identical on every single Microsoft domain that's created, so you don't want to mess around with them. Here's the computer half of Group Policy, here's the user half of Group Policy, and this file usually tells me not only my configuration settings, but what version of the file I'm using.

Creating the Central Store

3:54-4:21

To make a central store, all I do is--up here in the domain /SYSVOL /domain--I would create that PolicyDefinitions file. I could literally Copy my local PolicyDefinitions files right up here. It's just that easy. I have now created a central store.

Summary

4:22-4:49

Make sure, if nothing else, that you know the central store is for centrally managing administrative templates. You definitely need to know they're called .admx, or language specific, or .adml, and that I take my local policy definitions. I wouldn't worry too much about where it's located, but you have to know it's called PolicyDefinitions, and I'm going to make a copy of that up in the SYSVOL. If you've got that mastered, you're really going to be in good shape.

8.2.3 Group Policy Scope

Group Policy Scope
0:00-0:03
We're going to talk about controlling Group Policy scope.
Group Policy Processing Order
0:04-0:28
In another lesson in this class, we'll talk about the Group Policy processing order. In an ideal situation, that's the best way to control what settings affect the workstation. These Group Policies are processed in a particular order, whichever one runs last wins. There may be circumstances where that's not enough to make sure that you get the desired outcome.
Four Ways to Control the Scope
0:29-0:44
In that case, there are ways to control the scope, and really, there's four of them. One is to Block Inheritance, another is to Enforce a policy, we have Security Group Filtering, and we have WMI Filters.
Here, I've just drawn a little diagram.
Example
0:45-1:42
I've got a domain, Northsim.com. I've got an OU called Sales. Inside of Sales, I've got my Sales Desktops. I have another OU named IT. They have a sub-OU for Desktops as well, and then I've got a couple of policies applied to the domain.
One policy says No USB drives. Another policy installs an Anti Virus. Let's say, for example, that this No USB Drives policy is going to affect everybody in the domain, so that's going to be all the Sales desktops and all the IT desktops, but you want your IT desktops to accept USB drives. You actually don't care who's logged in, you're simply presuming that it's only IT people logging in to these desktops, and you do not want GPO1 to take effect on any computer that's inside of this OU.
Block Inheritance
1:43-3:15
In that case, we could block inheritance. When you block inheritance, it actually makes a little blue circle on the OU. There's a white exclamation mark in it, and we'll see that in the demo. You do need to be familiar with how it looks. Block inheritance basically says that any policies above that container will be blocked. It's really what Group Policy Inheritance is. It just means that I'm accumulating all the policies that run before wherever my account lives.
If I block inheritance at IT Desktops, none of the policies at the Site level and none of the policies at the Domain level are going to run. Only those policies directly attached to IT Desktops will be processed by the computers inside of that OU. Now, in real life, if this was my situation, I probably would just attach a policy here that says, "Yes, USB drives," and call it good, but we're blocking inheritance just to take a look at some of our options.
If we really do mean everything above a particular level should not run, you Block Inheritance. Now, we've got a problem, because I've got this GPO2 that's installing Anti Virus, and I don't want that to be blocked. I'm okay with everything that's being blocked from the Site, I'm okay with any other policies on the Domain, but this particular policy, I do not want it to be blocked or overwritten. In that case, I would set it to be enforced.
Enforced
3:16-4:16
When you set a policy to be enforced, it's going to put a little lock on the policy, which we'll see in the demo, and what that means is that policy cannot be blocked and it cannot be overwritten. Now that I have enforced the Anti Virus GPO, whether there's a block at Desktops, it's still going to break through that Block Inheritance. Or if I had a policy that uninstalled Anti Virus, for example, that would not take effect, because this cannot be overwritten by any other policy.
It gets a little messy if you enforce multiple policies. This is a last ditch thing. You want to just enforce the policies you're really sure they've always got to take effect. Now, we can use Block Inheritance and Enforced; those are very wide-filtering. I'm Blocking Inheritance, so I'm really blocking a whole lot of policies, and I'm doing it at the OU or the Domain or the Site level, Enforcing, same thing.
Security Group Filtering
4:17-4:30
A lot of times what we're interested in is having a policy only applied to a particular group of users or only not apply to a particular group of users, and they call that Security Group Filtering.

Negative Filtering

4:31-6:09

In this example, I've got my same domain, but now, the person who designed this domain put all the desktops in one OU, but we still have our No USB Drives GPO. We don't want that to take effect for any of the Admins in the IT OU. You can't Block Inheritance, because all my desktops are in the same OU. Security Group Filtering is used when people that it should apply to or not apply to are scattered throughout. I can't deal with it on an OU basis, because OUs are not Security Principles.

What I've got to do is this: I've got to make a group that groups together all of the people that I'm going to either apply this or exempt it. I'll make IT Admins, and I'll make a Global Group. In this case, I don't want my No USB drives GPO to apply to anybody in this group. I'm going to go into the Properties of this GPO, and I'll show you exactly where in the demo, and I'm going to add the IT Admins Group, and I'm going to give them two rights.

The first right I'm going to give them is to Deny Apply Group Policy. That right there is sufficient. At that point, they cannot apply that policy, therefore, they would be exempt from it. In practice, you should also deny them Read, because if you just Deny Apply Group Policy, the computer has to read the policy and then not apply it. If you Deny Read, it makes login a little bit faster for them. That will go ahead and filter the policy for everybody that's in that particular group.

Positive Filtering

6:10-7:15

We can also do security group filtering in a positive way. Maybe you have a policy that should only apply to a particular group. Here I have my domain Northsim, but now what I've done is gone ahead and created a GPO that says yes for USB drives. And maybe in my default domain GPO I have them denied, and then this one is going to override that; it's a little bit lower in the link order. I only want this policy to apply to the users in my IT OU. Again, I'm going to create a group.

At my group, IT Admins_G, I made global group. What I'm going to do is go into the properties of this GPO. By default, all GPOs allow authenticated users to read and apply the GPO. I'll pull authenticated users off, and then the only one I will add to that section that says "Who can apply this GPO?" will be IT Admins. Then they will be the only group that can actually run that particular Group Policy.

WMI Filtering

7:16-9:09

The last way I have of controlling the Scope of Group Policy is something called WMI Filtering. WMI stands for the Windows Management Interface. You don't have to know that, you just really have to know what it does. A WMI Filter runs a test. The only thing that's a little tricky with this is you have to script it out, but for almost anything you would want to script, you can go out on the Internet and find an example of somebody who's already done it.

We create a script that runs a test. The answer to that test is either yes or no. If the answer is yes, the policy gets applied. If the answer is no, the policy does not get applied. Now, when you're designing WMI filters, keep a couple of things in mind. Number one, that test can be as intricate as you need it to be, so we can say, hey, I'm only going to look at computers that were manufactured by Dell and have 4 GB of free space on the hard drive and are SATA drives.

If it matches all three of that criteria, then my answer is yes. The policy gets applied. If any one of those are no, the policy won't get applied. You can also use ors. I'm looking for anything manufactured by Dell or by Gateway, either one of those is yes, my answer for the WMI filter is yes, the policy will get applied.

All the filter can do is give you a yes or a no. Yes applies the policy; no does not apply the policy. It can't say if the answer is yes, I'll apply some policy and if the answer is no, I'll apply a different policy. That cannot be done. If that's your situation, you have to create two different policies and two different WMI filters and set it up that way.

WMI Examples

9:10-11:29

Here, I've created a couple of examples we can look at. We've got my NorthSim.com domain, we've got Sales, IT. Now we've put all of the computers in Corp Computers. There not just desktops, I've got my laptops in here as well. I have a GPO that's going to distribute some laptop drivers, but I only want it to apply to the laptops. I would create a WMI filter that says, "Is this computer a laptop?" If the answer is yes, it will get the policy. If the answer is no, it won't get the policy.

I have another GPO2 that's going to install Office, and let's say, this is going to install Office 2007. I want to make sure that before Office 2007 gets installed, a couple of different criteria are being met. Number one, I've got enough free space, because I know it's going to take 2-3 GB, and if the laptop or the desktop doesn't have that amount of free space, it could crash it. I'm going to run a test to make sure there's a certain amount of free space, and maybe I only want 2007 on XP. If the answer is yes for any of these computers--they're running XP, and they have 2 GB of free space on the hard drive--then Office 2007 will get installed.

If I had Office 2010, that was going to get installed on Windows 7 or Windows 8, I'd have to do another GPO with another WMI filter. So in that case, I would have to create another GPO: GPO3. That will install Office 2010, but only if I have my free space and it's Windows 8. We're controlling Group Policy for the most part, using the processing order. If we need to limit the Scope of the Group Policies, we have four ways to do that.

We can Block Inheritance, which will prevent any policies above that level from being run. We can Enforce the policy, which will make sure that it cannot be blocked, cannot be overwritten. We can use Security Group Filtering to say that the policy either only applies to a particular group or only does not apply to a particular group. Or we can use WMI Filters that run a test. The answer to the test is yes, we apply it; if the answer to the test is no, we don't apply it.

Summary

11:30-11:37

That will give us all the flexibility we need to make sure that every user and every computer gets exactly the configuration they need with Group Policy.

8.2.4 Configuring Group Policy Scope

Configuring Group Policy Scope
0:00-0:23
In this video, we're going to take a look at controlling the scope of Group Policy. I want to go into group policy Management Console, go Tools, Group Policy Management, and you can see I've got a number of policies here. Some are linked to the domain. Some are linked to some OUs. I don't have any linked to the site, but I could.

Group Policy Processing Order
0:24-0:58
Now we know that because of Group Policy Processing Order, first Site policies are processed, then Domain policies, then policies at the OU, and finally policies at the sub-OU if there's nested OUs. If there's multiple policies at the same level, I can go in and change the link order to control which one is processed first. This is probably the best way to manage Group Policy. If you can just manage it using the link order, that's your best bet.
Sometimes you can't. In that case, we're going to see what we can do to control that.

Block Inheritance
0:59-2:15
Let's just say management said "No, Shad, you're not allowed to have an Allow USB Drives on the Sales OU. So I'm going to delete my link. But up at the domain, I have a Block USB Policy, and I need for those people in Sales to be able to use their USB drives. Now if I need to block a policy, then I can turn on Block Inheritance.
If we take a look, right now, the users and computers at this level are first going to get the block USB drives, then the Default Domain Policy, and then Redirect Documents, because that's attached to the OU. When I Block Inheritance, essentially what I'm saying is any policy that runs at a level higher than this. So, all the Site Policies, all the Domain Policies are going to be blocked. If I right click it and I say Block Inheritance, you can see right away on the Inheritance tab, everything at the domain level went away.
Now this is not a great idea. You can see it even puts a little blue circle with a white exclamation point there just to let you know not a great idea. I'm not in favor of Block Inheritance, but I want to show you what it does. You should know what it does, and you need to be familiar with any of the icons that you would see in here, so even if they don't tell you it's Blocked Inheritance, you see that little exclamation mark. You know that's what the problem is.

Enforced
2:16-3:09
Now you might have a policy at a level and you say, "look, I need to make sure that this policy can never be blocked and never be overwritten." Well, in that case, I can set it to Enforced. When I set it to Enforced, it's going to break through that Block Inheritance. It also is going to override any policies that would try to run after it. This is the policy that's going to win, period. End of discussion.
We'll go ahead and we'll enforce the Default Domain Policy. Right click it and hit Enforced. You can see it gets a little lock right there. Click on Desktops. Click on Domain Controllers; you might be able to see it a little better. A lock means that it's Enforced, and if I go down to Sales and I look at my Inheritance, you can see it's brought back the Default Domain Policy, and it's jumped it to the top of the list. Regardless of what happens, this is going to be the policy that wins. I'm going to unenforce and unblock.

Security Group Filtering
3:10-3:36
A third method we have for controlling the scope of Group Policy is called Security Group Filtering. A lot of people get confused. They say, "Group Policy"-- it's a policy for groups. No, it's a group of policies. It lets us change settings on the computers based on the identity of the computer, or the user, or both. The only place where groups come into play for Group Policy is with Security Group Filtering.
There's two ways to do this.

Positive Filtering
3:37-4:37
One way is to make sure the policy only applies to that group. The other way is to make sure the policy applies to everybody except that group. We'll go ahead and take a look at both. Let's say we have a policy, Redirect Documents, but we only want this to apply to members of the Sales Users Group. Maybe there's other groups inside that OU. There's Sales Managers--I don't want it to apply to them. I just want it to apply to Sales Users. In here, you can see it says "Security Filtering". Right now, it's allowing all authenticated users to apply this policy. What I would do here is, add in my SalesUsers, and then get rid of Authenticated Users.
So now, in effect, this policy only applies to members of the SalesUsers group. You want to be aware if you have users all over the place, this is the only way that you can do this. If these SalesUsers were scattered throughout lots of different OUs, then I'd have to use Security Group Filtering.

Negative Filtering

4:38-6:08

Maybe I have a situation where I say, "yeah, but I need to exempt Sales Users from these Blocked USB drives." I can't take away Authenticated Users from Security Filtering, because I do want all Authenticated Users to get the policy, I just don't want it to apply to members of the SalesUsers group. In that case, that's still called Security Group Filtering, even though you won't see it in the software. We would go up to the Delegation tab and we would Add in SalesUsers. As you can see, it's actually just saying Read.

We need to go, and we need to add them in, and we need to deny them the ability to apply this policy. I'm going to come now to Advanced. You have to click Advanced in order to do this. I'm going to Add in SalesUsers. Now in effect, I just need to deny them Apply Group Policy. That right there will exempt them from this policy.

In real life, also go through and deny them Read. The reason is this: if you leave Read on but Apply Group Policy is off, the computer will read through the entire policy and then not apply it. Well, there's no point in the computer reading through that policy for these users; it's not going to apply to them. By milliseconds, you can slightly speed up the processing by denying them Read. You're setting it Denying Permission; it's going to take precedence over Allowed-- do you really know what you're doing? Absolutely. Now I've used Security Group Filtering to effectively exempt members of the SalesUsers from this policy.

WMI Filter

6:09-8:52

The last way that I can control my Group Policy Processing is using WMI filters. WMI filter runs a test. If the answer is yes, it applies the policy. If the answer is no, the policy does not get applied. This is a little tricky, because we're starting to get into scripting here. Let's go ahead and make a New filter, and I'm going to call this "Space Free on C Drive" and then you need to add in the code that you're going to use.

What this does is it goes out and it gets all the information from the Win32 Logical Disc. Then it says it's looking for where free space is greater than--and that number is roughly 2 GB so it's running a test to make sure I have at least 2 GB free on the hard drive. If the answer is yes, the policy gets applied. If the answer is no, the policy does not get applied.

Once I've got my WMI filter, then I can apply it to my policies. Maybe I have a policy that's going to install Microsoft Office on the computer, but I only want it to do that if it has sufficient free space. Let me go up to Scope, and down here I can set a WMI filter. I can only set one WMI filter per GPO link, and again, it's a yes/no. I can't have it test and say, "well if it's Windows XP, do this; if it's not, do something else. It's just if the answer is yes it runs the policy. If the answer is no, it does not run the policy.

If we want to speed up policies a little bit, we can come over to Details-- you can see that my policy is enabled. If you know you've only set something up in the computer side of the policy, you could disable the user side. Or, if you've only used the user side, you could disable the computer side. That way, any users within the scope of this policy know there's no settings on the user side. They're not even going to bother to look at this particular policy.

I can also go up under Settings and I can see what Settings have been defined. We'll take a look at the settings for the Default Domain Policy, and under Settings I can see what settings have been defined and see exactly what that policy does. Now if I've got policies overriding each other, I can just look at the Inheritance tab. I can see what policies are going to win.

On Delegation I'll see which users have permissions for the OU, but it's not going to tell me anything about the result of Security Group Filtering, Blocking Inheritance, Enforcing Inheritance--it'll tell me a little bit on that, but the point is I can't see exactly what my effective policy is going to be. If you are controlling the scope, you want to know about Group Policy Modeling and Results.

Group Policy Results

8:53-9:21

Results just gives me a quick and dirty result for a user and a computer combination. I right click Group Policy Results Wizard. Next, this computer, this user, go. I can see on the Details exactly what my settings are going to be and which policy is driving that setting. If you have a situation where somebody can't do something, and you think, "well, they should be able to do that," you can come in here and do Group Policy Results and see exactly where it's coming from.

Group Policy Modeling

9:22-10:21

If you want to test what will happen if you implement a policy, that would be Group Policy Modeling. I right click Group Policy Modeling Wizard. First of all, I can actually see what would happen if it's processed by a particular domain controller. If I'm concerned that not all the policies have replicated, I can choose a domain controller. I can show domain controllers from other domains if I have a forest with more than one domain.

What's going to happen if a domain controller in another domain authenticates this user? I can set up users from a particular container or a specific user. Same for computers. Computers from a specific container or a specific computer. We'll say, "what happens when HR Users log on to Sales Computers?" Now, I can go through and simulate a slow network connection automatically by default. If the connection falls below 512 kbps, there are certain parts of the policy that would not be processed, like software install--things where slow network connection would affect it.

Loopback Processing

10:22-11:41

I also can set up Loopback processing. Loopback processing is used for kiosk machines. Basically what it says is this: if we set it to Merge by default, the computer half of a policy runs before the user half. There aren't many settings that are on both sides, but by default, the user half would win because the user half runs when the user logs in, which is after the computer boots.

With Merge, what it's saying is if there's a conflict between the user half and computer half, now the computer half is going to win. If I say Replace, it's actually going to ignore user policies altogether and just run the computer side of all the policies that are in effect. I can set up a particular Site. I want to simulate them logging in from a Site--see what Site policies will affect them. I can simulate adding the user to a particular group. I can simulate adding the computer to a particular group. I'm testing to see if Security Group Filtering is going to affect them. I add them to that group. I can test if I have WMI filters that apply to that user, WMI filters that apply to that computer, and then I'm good to go.

Anything that will affect the outcome of all these events going together is going to be processed, and then I'm going to see the details. I can see exactly what the effective settings are going to be and then where they're coming from. Group Policy Modeling and Results Wizards are great for troubleshooting Group Policy.

Summary

11:42-12:05

Controlling the scope of Group Policy, we can Block Inheritance, we can enforce policies so they can't be blocked--they can't be overwritten. We can use Security Group Filtering to make sure the policy only applies to a particular group, or make sure a particular group is exempt from that policy. We can use WMI filters to run a test. If the answer is yes, they get the policy. If the answer is no, they don't get the policy. That's how we control the scope of the Group Policy objects.

8.2.5 Group Policy Management Facts

A Group Policy object (GPO) is a collection of settings that can be applied to a group of users or computers. A number of factors determine the effective Group Policy settings for an object. When working with Group Policy settings, be aware that:

- Through Group Policy inheritance, settings in a GPO are applied to all objects below the container where the GPO is linked. Inherited GPO settings for any object are the total settings of all GPOs linked to all parent objects.
- GPOs are applied in the following order:
    1. The local Group Policy on the computer.
    2. GPOs associated with a site.
    3. GPOs linked to the domain.
    4. GPOs linked to the organizational unit (OU). If the OU has nested OUs, the Group Policy is applied from the highest-level OU to the lowest-level OU. In other words, the Group Policy in the parent OU will run before the Group Policy in the child OU.
- A specific setting in a GPO can be:
    - Undefined, meaning that the GPO has no value for that setting and does not change the current setting.
    - Defined, meaning that the GPO identifies a value to enforce.

    Be aware of negatives in policies. If you disable a policy that disables a feature, the feature is enabled.

- Individual settings within all GPOs are combined to form the effective Group Policy setting as follows:
    - If a setting is defined in one GPO and undefined in another, the defined setting will be enforced (regardless of the position of the GPO in the application order).
    - If a setting is configured in two GPOs, the setting in the last applied GPO will be used.

    The Local Group Policy is applied only when there are no GPOs linked to a domain or the OU. GPOs linked to an OU override GPOs linked to a domain when both are applied.

Scoping is the process of targeting a GPO to specific users and/or computers. Scoping methods are listed in the following table:

Method	Description
Block Inheritance	Blocking inheritance prevents settings in all GPOs linked to parent objects from being applied to child objects. • You configure inheritance blocking on the domain or an organizational unit (OU). • You cannot block inheritance on a per-GPO basis; blocking inheritance blocks all GPOs linked above the blocking object. • Only Group Policies applied directly to the container take effect. • A blue circle with a white exclamation mark in it indicates blocked Group Policy inheritance.

Enforced	To prevent inheritance from being blocked for a specific GPO, select the **Enforced** (no override) option for the GPO link. • You configure the enforced option on a per-GPO basis. • Enforced GPOs are applied last and override other GPO settings. • An enforced policy cannot be blocked or overwritten. • A lock icon indicates an enforced policy.
Security group filtering	To use Security group filtering: • Create a global group. • Filter in one of the two following ways: Filter a policy you want to apply to everyone but the global group by setting the following rights for the global group: ▪ Deny - Apply Group Policy ▪ Deny - Read Filter a policy you want applied only to the global group by modifying the properties of the GPO to allow only the global group to run the Group Policy.
Windows Management Interface (WMI) filtering	Use Windows Management Interface (WMI) filtering to determine the scope of a GPO dynamically, based on hardware and software characteristics such as CPU, memory, disk space, registry data, drivers, network configuration, or application data. In WMI filtering you create a script containing a test that results in a *yes* or *no* response. WMI filtering: • Applies the policy if the response is *yes*. • Does not apply the policy if the response is *no*. • Is restricted to only one WMI filter per GPO. • Uses queries written in WMI query language (WQL). • Should be applied for a well-defined purpose and limited amount of time. • Evaluates the target computer every time a Group Policy refresh occurs.
Loopback Processing	By default, Group Policy configuration applies computer settings during startup and user settings during logon. For this reason, user settings take precedence in the event of a conflict. With *loopback processing*, computer settings are reapplied after user logon. Following are some circumstances when you might use loopback processing: • If you want computer settings to take precedence over user settings. • If you want to prevent user settings from being applied.

	• If you want to apply specified user settings for the computer, regardless of the location of the user account in Active Directory. Loopback processing runs in **Merge** or **Replace** mode. • **Merge** mode gathers the Computer Configuration GPOs and appends them to the User Configuration GPOs when the user logs on. • **Replace** mode prevents the User Configuration GPOs from being applied.

For each GPO, the following options in Group Policy Management help you to manage the application of the GPO:

- On the **Details** tab, set the GPO Status to reflect how the policy is applied:
  Use the **Computer configuration settings disabled** setting if the Group Policy applies only to users or groups.
  Use the **User configuration settings disabled** if the Group Policy applies only to computers.
- On the **Settings** tab, you can view the settings that have been defined.

To determine how scoping affects the application of the GPO:

- Use **Group Policy Modeling** to launch the Group Policy Modeling Wizard: You can simulate how the Group Policies will be applied:
  Based on a specified user or users in a container.
  Based on a specified computer or computers in a container.
  Based on a slow network connection.
  Based on Loopback processing.
- Use **Group Policy Results** to launch the Group Policy Results Wizard and determine how Group Policies are applied for a specified user and computer combination. The **Details** tab of the Group Policy Results Wizard identifies settings as well as the Group Policy driving each setting.

8.3 Password Policies

As you study this section, answer the following questions:

- In what ways can password policies be set up?
- What strategies should be implemented to protect against password attacks?
- Which object types can be associated with a granular password policy?
- What are the characteristics of a strong password?

After finishing this section, you should be able to complete the following task:

- Configure a password policy and apply it to specific users or groups.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Manage Security Policies
  - Configure Account Password Policies

8.3.1 Password Policies

Password Policies
0:00-0:10
Let's talk about password polices. Password polices are very important in an organization, because you want to make sure that your users have good password hygiene, so for example, they have passwords that are long enough that they would defeat a brute force attack--somebody just trying different combinations of numbers and letters--but not so long that the user can't memorize them.

Password Length
0:11-0:35
If you're going to say that they need a 30-character password, everybody and their brother is going to write that down on a sticky note. It's probably under their keyboard, or on the side of their cubicle, or the side of their monitor. The password should be long enough to be secure, not so long that it can't be remembered.

Change Password Often
0:36-0:42
We want to force them to change the password every so often, usually somewhere between 30 to 45 days.

Complex Passwords
0:43-0:55
We also want to go through and make sure that they use complex passwords, which means a combination of three out of the four, meaning upper case, lower case, numbers, and special characters, so the three of those at least being present in the password.

Minimum Amount of Time To Keep Password
0:56-1:25
Along with the length of the password and how often they should change it, we're also going to go through and set the minimum amount of time that they need to keep that password. A lot of people have a favorite password-- their kid's name, or their dog's name-- and they figure out Shad makes me have five unique passwords, after the fifth password, I can come back to my favorite one. So they'll just change it five times all at once, and they'll get back to the one password that they've been using the whole time.
We'll go through and set up all these password options.

Account Lockout Policies
1:26-1:55
To go along with your password policy, you set up account lockout policies. When I sit down at the computer, how many times can I type my password wrong before I get locked out? How much time goes by if I typed it wrong twice and I have two strikes against me? How much time does it take before those are erased, and if in fact, if I do get locked out, how long am I locked out for?
It's not so critical to have memorized the exact parts of the password policy, but what is important to know is this: in Windows Server 2003 and before, the only password policy that took effect on the users was the one that was set up in the Default Domain Policy.

Password Policies In Windows Server 2003
1:56-2:56
If you added a different Group Policy--let's say to an OU--and set up a different password policy, the only way it would affect users is if they are local users on that machine. The domain users were all limited to whatever was set up in that Default Domain Policy. That was kind of a problem, and let me give you an example: if I have a company and I say, I'm going to have my employees have eight character passwords, and they'll change them every 30 days. Then the military comes to me and says, "Shad, we love your company and we want to give you a really expensive contract, but anybody who works on our contract is going to need 40-character passwords." I'm exaggerating a little bit. You get the idea.
I don't want to inflict that on all my users, so with Server 2003 and before, I'd actually literally have to create another domain so that I could have a different password policy.

Fine Grained Password Policy or Password Setting Object (PSO)

2:57-4:08

Starting with Windows Server 2008, Microsoft implemented a new type of password policy called a Fine-Grained Password Policy. It's also known as a Password Settings Object and in the demo, we'll take a look at why those words are interchangeable.

Either one--Fine-Grained Password Policy or Password Settings Object PSO--either of those are the same thing. What they are is a password policy that can be applied to a user or a group that would give them different password and account lockout settings than are defined in the Default Domain Policy. Anytime I need a particular user or group to have different password settings, I'm going to create a Password Settings Object.

A couple things to note about this: once you create your Password Settings Objects, you're going to have to go through and specify to whom this Password Settings Object applies, so first I create the PSO, then I add the user or group to that PSO, that's the first thing. The second thing is these Password Settings Objects or Fine-Grained Password Polices can only be applied to users or groups. They cannot be applied to OUs.

Shadow Group

4:09-4:57

Let's say you do have an OU, the ITOU, you say, "Well, my regular users want to have eight character passwords, and I want my LAN Administrators to have a little bit longer, they're a little bit more secure, because they have more authority within my network," so I'll make them have 12 character passwords.

Even though all my Domain Administrator accounts are in my ITOU, I can't apply the ITOU to the Password Settings Object, so what I would need to do is create what they call a shadow group. A shadow group is just a group whose membership is identical to the users that are in that OU, so I would create a shadow group IT membership there. Every user that's in the ITOU will be a member of the IT group. Now I've got my shadow group, I create my Fine-Grained Password Policy, and then I apply that group to the policy.

Summary

4:58-5:16

Password Policy is very important, it specifies what makes for good password hygiene in our network, and what happens when you don't remember what your password is, Account Lockout Settings, and then if we need anybody to have different settings than the default domain policy, we'll create a Fine-Grained Password Policy, also known as a PSO.

8.3.2 Configuring Password Policies

Configuring Password Policies
0:00-0:19
In this video, we're going to take a look at creating Fine-Grained Password Policies otherwise known as PSOs. I'm going to show you two ways to do it. It's a lot better in Windows Server 2012, but I do want to show you the old way in case it comes up anywhere. I do want to show you the old way in case you ever have to do it in real life on an older machine.

Password Policy in Windows Server 2008 and 2008 R2
0:20-0:27
In Windows Server 2008 and 2008 R2, the utility that we would use to create PSOs is ADSI Edit.

PSO
0:28-1:05
What is PSO? It's a password policy that we can apply to specific users or groups. Let's take a look at the default password policy, and then we'll focus on creating our new Fine-Grained Password Policy. By default, in a domain, the only password policy that's in effect for the users is the Default Domain Policy. If I set a GPO at the OUs with a different password policy, it only applies to local accounts on the computers. It does not apply to any domain accounts, so domain accounts are only subject to whatever's been set in the password policy of this specific Group Policy.

Default Password Policies
1:06-1:20
We'll take a look at the defaults. I'm going to open up Policies, Windows Settings, Security Settings, Account Policies, and here's my Password Policy. Enforce password history, this is going to remember 24 passwords, so it's going to take them 24 passwords to get back to their favorite password.

Enforce Password History
1:21-1:27

Minimum/Maximum Password Age
1:28-2:12
They need to change their password at least every 42 days, but they need to keep the new password at least one day. It's actually important. A lot of users have a favorite password. Let's say my favorite password is Shad1234. If you make me change it, I'm just going to change it to Shad2345, and then Shad3456, and I'll keep going. If I figure out that you're only remembering five or six passwords, I could actually just sit there, in five minutes change my password five times, and get back to my favorite one. By forcing me to keep my password for a minimum amount of time, you prevent me from just changing it a whole bunch of times right in a row to get back to the password that I really want.

Minimum Password Length
2:13-2:26
Minimum password length is seven characters. Usually, you want to use at least eight. I wouldn't go much further than 10 or 12, because at that point, people really can't remember them, and then they start writing them down, which is just as big of a security risk, if not more.

Password Must Meet the Complexity Requirements
2:27-2:38
"Password must meet the complexity requirements", means that the password must have three out of the four characteristics; upper case, lower case, numbers, and characters. So they need to use at least three of those.

Store Passwords Using Reversible Encryption
2:39-2:52
"Store passwords using reversible encryption"--that doesn't even sound good. Reversible Encryption was used to store passwords back in Windows NT 4.0. It would just be used for backwards compatibility with some type of a legacy application.

Account Lockout policy
2:53-3:44
Part of my Fine-Grained Password Policy will be the password policy, but will also have a section where we can specify the account lockout policy. Lockout threshold is how many times I can hack my password before I get locked out. "Lockout duration" is once I get locked out how long will I stay locked out. If you set that to zero, it means indefinitely. I come in this morning, let's say I type my password wrong a couple of times, and then I get logged in. "Reset account lockout counter after" is how much time needs to go by before those two strikes against my account are removed. If you have a very dynamic environment you'd keep that short-- very stable environment, users who don't often forget their passwords, you could keep it a little longer. These are the settings we're going to be setting up in our Fine-Grained Password Policy.

Configuring Password Policies Using ADSI Edit
3:45-4:09
First, let's take a look at ADSI Edit. The first thing I do is connect up to my computer. We can connect up to the Default naming context, and I can see in here my domain.

Password Settings Container
4:10-4:20
I'm going to open up System, and in there, I'll find my Passwords Settings Container. So we're really looking at the raw Active Directory. Now, hopefully, you can see where those distinguished names come from.

Make a New Object
4:21-4:57
Once I get to my Password Settings Container, I'm going to right click and make a New Object. I am creating a Password Settings Object. That's where the term Password Settings Object or PSO comes from. Consistently in tests, you will see it referred it to either way as a Fine-Grained Password Policy, or as a Password Settings Object, or even just as a PSO. I'm going to hit Next. This is the name of my PSO.
The first thing I need to specify is the Precedence.

Precedence
4:58-5:46
This is an arbitrary number that I assign. What it's used for is this: if there's a PSO applied specifically to that user, that's the PSO that's going to be in effect. However, if there's not one applied to the particular user, and the user belongs to two different groups that both have Fine-Grained Password Policies, this Precedence is going to decide which policy will take effect. If I'm a member of two different groups-- one has a PSO with a precedence of one, the other one has a PSO with a precedence of 10--the PSO with the precedence of one is the one that's going to affect me. You can set whatever number you want, just make sure the PSOs that are more important get a lower number, the ones that are less important get a higher number. Hopefully, you're not going to be creating a ton of these things.

Password Policies
5:47-7:14
Now it wants to know if I'm going to store my password with Reversible Encryption. Now False, you don't want Reversible Encryption unless we have NT 4.0 in the environment. It wants to know the Password History Length. How many passwords are we going to remember? The domain remembers 24, so we'll go ahead and remember 24 as well. Am I going to require them to have a complex password? Absolutely, that is true. What will be their Minimum Password Length? For these Domain Admins, I'm going to make them have 10 characters at least. Now we have the Minimum Password Age, that's the minimum amount of time I must keep my new password before I'm allowed to change it. We'll go ahead and set that to three days. Now when you're working in ADSI Edit, it's a little weird because you have to specify days, hours, minutes, and seconds, even though you might not be working with all of those fields. You'll see if I try to click Next and when I get to the end it'll yell at me. You can't do it this way, so I'm going to go back and actually set it up correctly. Three days, zero hours, zero minutes, zero seconds. Maximum Password Age, maximum amount of time that they can keep their password. The domain was 42, but we're going to set this back to 30, so 30 days, zero hours, zero minutes, zero seconds. Account Lockout Threshold (how many times can they hack their accounts before they get locked out).

Account Lockout Policies
7:15-8:11
We'll let them try three times. These are admins, they should remember their password. Lockout Observation Window is when I come in, I type my password in wrong twice, I've two strikes against me, how long goes by before those two strikes are removed? This should be shorter than the amount of time that they actually get locked out. We'll give them maybe an hour, zero days, one hour, zero minutes, zero seconds. Now how long do they get locked out? Let's keep them locked out for an hour, zero days, one hour, zero minute, zero seconds. Then, I can go ahead and click Finish, and now I have my PSO.

PSO Applies To
8:12-8:51
At this point, the PSO does not apply to anyone. As a separate second step, I need to come in here, scroll down, and find PSO Applies To, and then I would go in and add a Windows account. Again, it can either be users or groups to which this is going to apply. Now this new PSO that I've created is going to apply to the Domain Admins Group. This is not the most friendly way to create PSOs.

Configuring Password Policies Using Active Directory Administrative Center
8:52-9:10
Luckily, with Windows Server 2012, they came out with a little bit better way. I can go ahead and use the Active Directory Administrative Center. I would go into my domain, and just like I did in ADSI Edit, I want to open up System.

Password Settings Container
9:11-9:34
Inside of System, you're going to see the Password Settings Container. Either double click it, or you should be able to get it from here, and click New Password Settings. You can see this is a lot more friendly. Just like in ADSI Edit, I have to set my Precedence.

Precedence
9:35-9:41
I can choose to Enforce a minimum password length if I want to.

Password Policies
9:42-10:05
This is seven characters, we'll give them nine. I can choose to Enforce my password history. It's set to the domain default of 24. Password must meet complexity requirements. There's my reversible encryption. Here's my minimum password age, maximum password age, and then I can go through and adjust my account lockout if I need to.

Account Lockout Policies
10:06-10:38
Number of failed logon attempts--three. Reset failed logon attempts after--this is showing 30 minutes. "Account will be locked out for"--that number, again, needs to be longer than this one right now, or the same, or I can just say indefinitely. Same as putting zero inside the password policy. I also can choose whether or not to protect this PSO from accidental deletion.

Directly Applies To
10:39-11:31
Remember, it's always a second separate step to specify to whom this policy applies. I still need to come down in here and add some type of user or group to whom this password policy is going to apply. With groups, it's only going to let me choose from Global Groups. You can see if I do a Find Now I see Global Groups, I also will see any user accounts that exist in my particular domain. We'll apply this to domain guests. And I've got my group, or it could be a specific user down there. If I look in my Password Settings Container, I see both of those policies. This is a great graphical way to see the policies and to see the precedence.

Configuring Password Policies Using PowerShell
11:32-11:36
The last way you could create a PSO is through PowerShell.

New-ADFineGrainedPasswordPolicy Command
11:37-11:50
To create the Fine-Grained Password Policy, it's just New-ADFineGrainedPasswordPolicy. If I hit Enter, it's just going to prompt me for all the same parameters we just put in.

Add-ADFineGrainedPasswordPolicySubject Command

11:51-12:23

Once you've created your Fine-Grained Password Policy, again, a second separate step is to add somebody to the policy. It's a little bit different. It's Add-ADFineGrainedPasswordPolicySubject. If you're memorizing things for tests, just remember these are Fine-Grained passwords--New, with the word FineGrained in the other half, makes a new policy. As a separate step, I add my subject. So it's going to start with Add, have FineGrained in it and end with Subject. That's how we handle passwords within the domain.

8.3.3 Password Policy Facts

Password policies define characteristics of passwords that are enforced by the system, such as the minimum number of characters in a password or how often the passwords must be changed. With Windows Server 2008 and later, there are two ways of setting password policies:

Method	Description
Account policies	Account policies control passwords and login properties for the entire domain. • Password Policy settings control characteristics enforced for user passwords. Account Lockout Policy settings control what happens when a user enters one (or more) incorrect passwords. • Settings in the local GPO are used if the computer is a member of a workgroup. Settings in the domain GPO are used for computers that are members of a domain. • Policy settings are applied to the computer, not the user. • Although you can configure Account Policies settings in any GPO, only the settings configured in a GPO linked to the domain take effect. The following list describes the password policy settings: • **Enforce password history** requires users to input unique passwords. Set this to a high number to keep users from frequently repeating passwords. Windows can remember up to 24 old passwords. A maximum password age must be configured for this setting to take effect. • **Maximum password age** requires the user to change the password after a given length of time. Setting this value to 0 means that the password never expires. • **Minimum password age** keeps users from changing passwords immediately after they've reset their passwords. This prevents users from defying the password history by entering several passwords to get back to a preferred password. The value must be less than the maximum age, and should be a setting greater than 0. A setting of 0 allows the user to reset the password immediately. • **Minimum password length** prevents users from using passwords that are too short. At a minimum, enforce passwords of 8 characters or longer. • **Password must meet complexity requirements** prevents using passwords that are easy to guess or easy to crack. This setting enforces the following: Requires users to create a password with a minimum of three of the four types of special characters (e.g., lower case letters, upper case letters, numbers, or !, @, #, $, %, ^, &, *). Disallows the use of dictionary words or any part of the user login identification. Requires that passwords are 6 characters or longer.

	• **Store passwords using reversible encryption** is essentially equivalent to storing plain-text passwords. This setting should be disabled unless a specific application requires access to the plain-text password. The following list describes account lockout policy settings. • **Account lockout duration** determines the length of time the account will be disabled (in minutes). When the time period expires, the account will be unlocked automatically. When set to 0, an administrator must unlock the account. • **Account lockout threshold** determines the number of attempts a user can make before the account is locked. A typical setting is 3. • **Reset account lockout counter after** determines the amount of time (in minutes) that must pass before the number of invalid attempts counter is reset.
Granular password policy	Granular password policies allow you to create password policies for users and global groups separate from the password policy applied to the entire domain. Using granular password policies, you could, for example, require administrators to use 14-character passwords, while requiring only seven-character passwords from standard users. You should know the following facts about granular password policies: • The domain must be running at the Windows Server 2008 domain functional level or higher. • Password policies affect only user account passwords, not computer account passwords. • Only members of the Domain Admins group can set granular password policies, but you can delegate the permission. • Granular password policies are saved as a Password Settings Object (PSO) in the Password Settings Container (PSC). There is one default PSC. It cannot be renamed, deleted, or moved. You can create additional PSCs, but they will not take effect. The PSC holds one or more PSOs. You can define multiple PSOs, each with unique password policy settings. • PSOs have attributes for all of the settings that can be defined in the Default Domain Policy, except Kerberos settings. • Policies can be applied to user accounts or global security groups. Each granular policy can be applied to multiple users and/or groups. Granular password policies affect only users within the current domain. • Policies are not enforced when applied to OUs, the domain, or other group types. To apply a granular policy to all users within an OU, create a global security group that contains all OU members. Apply the policy to the group.

	When you move a user account to a different OU, remember to also change the group membership so that the granular password policy no longer applies.

In general, use Account Policies to enforce a domain-wide password policy. Use granular password policies to enforce policies for groups of users that have more or less restrictive password policy needs than the domain-wide password policy.

Use the following strategies to protect against password attacks:

- Educate users on how to create and remember strong passwords. Enforcing strict password restrictions might actually weaken network security if you do not educate users about proper procedures to take to protect logon credentials. If users do not understand the restrictions that have been implemented, they might try to circumvent these restrictions by writing down passwords. Take the following measures to educate users:
  - Tell users that they should not write down passwords or share logon credentials with other users.
  - Teach users how to construct and remember complex passwords. For example, for the password **bw2Fs3d**, users might create the following sentence: *bob went 2 the "capital" Florist shop 3 times daily*.
  - Educate users about social engineering tactics. Instruct them not to respond to requests for passwords from administrators or other seemingly trusted personnel. Implement policies that prevent administrators from asking for sensitive information.
- Protect access to the password file. Passwords are typically stored in a password database file that uses a one-way encryption algorithm (hashing). Use methods available in the operating system to protect the password file.
- *Salt* the hash to mitigate rainbow table attacks. Salting the hash adds random bits to the password before hashing takes place, thereby producing an entirely different hash value for the password. Because the hacker does not know the extra random bits, the rainbow table will now be of no value.
- Implement two-factor authentication.

Password policies detail the requirements for passwords for the organization. This can include the following:

- The same password should never be used for different systems.
- Accounts should be disabled or locked out after a specified amount of failed login attempts.
- Passwords should never contain words, slang, or acronyms.
- Users should be required to change their passwords within a certain time frame and to use a rotation policy.
- A strong password policy should be enforced. Strong passwords:
  - Contain multiple character types, uppercase, lowercase, numbers, and symbols.
  - Are a minimum length of eight characters or more.
  - Use no part of a username or e-mail address.

8.4 Audit Policies

As you study this section, answer the following questions:

- When does the *account logon auditing* generate an event on the system for a local user account or for a domain user account?
- What is the difference between Account Logon auditing and Logon auditing?
- What is the difference between auditing for success and auditing for failure?
- What types of items does *object access auditing* track access to?
- How can you use auditing to track changes to Active Directory objects?
- When would you use *process tracking*?

After finishing this section, you should be able to complete the following task:

- Use Group Policy to enforce auditing and secure audit logs.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Manage Security Policies
  - Configure Audit Policies

This section covers the following 70-410 exam objective:

- 602 Configure security policies.
  This objective may include but is not limited to:
  Configure Audit Policy

8.4.1 Audit Policies

Audit Policies
0:00-0:31
Let's talk about audit polices. Your job as the domain administrator or the network administrator is to take away as many rights as possible from the end user. Remember, it is their job to cry and get them back. In an ideal situation, they should end up with exactly the amount of rights and permissions that they need to do their job. No more, no less.
Now the question immediately comes up, "Have you done that?" How do I know that I've gone in and I've given them enough rights to do their job, but not more than they need? And that's done through auditing.

Pros and Cons of Auditing
0:32-1:14
The great thing about auditing is it lets you make sure that the security policies are set up correctly. It also will help you if you need to investigate somebody who's doing something that they shouldn't be doing.
The bad thing about auditing is that it requires resources from the server. If you think about everything you do during the day, now imagine you have to write it all down in a notepad. Well certainly that's a lot more work, and it's going to be the same thing with auditing. When we turn on auditing, we want to make sure that we only audit those things that are important to us. So maybe we just need to know when people log on successfully, or we just need to know when people try to log on and they're not successful. Maybe we don't need both of those, or maybe we do need both of those.

Object Access
1:15-2:15
The only one that's a little tricky with auditing is Object Access, and that's literally what it sounds like--we're auditing people who are accessing some kind of an object. If you ever want to know what something does in Microsoft, just take the end word, move it to the beginning, and that's what it does. An object could be a number of things, it could be a file, it could be a folder, it could be a server. It doesn't matter what it is. If it's an object that you want to audit the access to it, it's always a two-step procedure, and that's where it's different. All the other auditing, you simply turn it on and it happens.
With Object Access, it's always two steps, because if you think about it, you really want to turn on Audit Object Access and then have events for every single file that's used on the computer recorded in the event log. Windows uses thousands of files just as part of the operating system. With Audit Object Access you turn on the auditing setting, then you need to go to the object, go into the Auditing tab, and define what types of events you're going to audit.

Configuring Audit Policies
0:00-0:48
In this video, we're going to take a look at auditing. Now remember, as a Network Administrator, it is your job to take away as many rights from the user as you possibly can, but it's still their job to cry and get them back. Technically, we call that the Principal of Least Security, so they should have the least amount of security necessary to do their job--no more, no less.
How do you know you've properly set up security? There's a lot of places that you can set up security. We have Permissions to objects, we have Rights inside the file system, we have Security Options--a lot of different places that you're going to go through and configure the computer. Auditing allows you to check to make sure that's been done correctly. It also allows you to try to troubleshoot if there's a situation where you think that somebody's doing something wrong.

Need for a Company Security Policy
0:49-1:17
Now, let's be realistic. I've been involved in multiple investigations at different companies where people come in and say, "Hey, Shad is doing XYZ, I'm making a complaint about it," and then you've got to investigate in coordination with HR to find out if that exactly is the truth and what's going on.
Make sure your company has a Security Policy, because if you don't have a policy that they sign that says, "I understand this is company property, not mine," then you may not be able to enforce the results of that investigation.

Group Policies Involved with Auditing
1:18-1:54
Let's take a look at the Group Policies that are involved with auditing. Since we're just interested in the settings, I'm just going to edit the Default Domain Policy, but in real life, you want to make sure you link a policy to the appropriate Scope. Now, traditionally, we'll find the regular auditing policies under Policies, Window Settings, Security Settings, Local Policies, and then in here, we see the Audit Policy. This type of auditing has been around for quite a while. Let me point out a few of the ones that we often see come up in real life or in exams.

Account Logon Events vs. Logon Events
1:55-2:34
First of all, an easy one to confuse is Account Logon Events versus Logon Events. A Logon Event occurs wherever I sit down at the keyboard and hit Ctrl + Alt + Delete. An Account Logon Event occurs wherever I'm authenticated. If I'm on a standalone machine, and I log on, I'm going to have one each of those events, because I'm using a local account.
In a domain, the Logon Events will occur at the client. The Account Logon Events will occur at the domain controllers. If there are multiple domain controllers in a Site, you actually have no idea which domain controller is going to authenticate the user. What you're going to need to do is look at the security logs on all of the domain controllers.

Event Log Subscriptions
2:35-2:50
There's a great new feature that came out with, I believe it was Windows Server 2008 R2, where you can create Event Log Subscriptions so that you can centralize the event log. That's something you'd want to check in if you're really trying to track Account Logon Events.

Audit Policy Change

2:51-4:06

Another one to keep in mind is Audit Policy Change. It becomes very difficult when it's the administrator that's the bad guy. You might go in and say, well, we found out that Shad, the Network Admin, is breaking into a particular file or folder and keeps writing that the boss is funny-looking with big ears. We want to catch Shad in the act, so we're going to go through and we're going to Audit Object Access, which means we're going to audit this person's access to objects. I'm an administrator. I can see that, so here's what I'll do. I'll come in, turn off Audit Object, break into the folder, write, "The boss is funny-looking with big ears," and then come back in and turn the policy back on.

To catch me, you want to enable Audit Policy Change, and basically, what this does is, it will record something to the Security Log any time people change a policy for User Rights Assignment, the Audit Policy itself, or a Trust Policy. Now, when I'm turning the Audit Policy off, that's going to get recorded to the Event Log. Again, I'm an administrator. I'm pretty savvy with these things, so I might just go in and empty the log, and sometimes, that's the best you can do is find out that the administrator emptied the log and there's a gap, and that's what you're going to use as your evidence.

Audit Policy Change, if you are investigating an administrator, that's something that you could turn on.

Audit Object Access

4:07-4:40

Audit Object Access is exactly that. It lets me audit access to objects. Now, this is one you want to be very careful with. Auditing does put extra burden on the computer. Think about what you do in your day, and now, think about doing all of that, except writing everything down in a notepad.

When we're Auditing Object Access, we're not actually interested in every single object in the computer. Objects are files, folders, printers. That would be way too much information. It's important to understand that Audit Object Access is always a two-step procedure.

Success and Failure

4:41-5:39

We go in and we define what we're interested in. If I am breaking into the folder and writing, "The boss is funny-looking with big ears," we're interested in Successful Object Access. If the complaint is somebody is attempting to break in and they haven't really got in yet, that's a Failure.

It might not be malicious. It might be a situation where an employee claims, "I can never get into that folder to save my reports up there." You're pretty sure that they can get in, but you want to find out what's going on, so you audit the failure to access that particular folder so that you can see when they're trying to save the file, how they're coming in, that sort of thing. It doesn't have to be an investigation but certainly, that's some of the context where you'll see this.

I would turn on my Object Access. Then, the second step is to go to the object itself and define what's going to be audited. Again, there's way too many objects in the computer for us to audit everything. We'll go take a look at that in a minute.

Advanced Audit Policy

5:40-6:43

I want to show you the Advanced Audit Policy.

In Windows Server 2008, they came out with an Advanced Audit Policy, but it was not available in Group Policy. It was done from the Command line using a command called auditpol. In 2008 R2, they integrated it into Group Policy so I can see in here. I've got an Advanced Audit Policy Configuration, and then underneath that, I can go ahead and click Audit Policies, and we've got quite a number of things that we can audit.

Most of the things I'm probably not even interested in half the time, but, I can see not only logons, but different logon events, credential validation. I can go through. I can set Detailed Tracking, Account Management, so I want to see when people are changing security groups, user accounts.

Object Access has been improved as well. Instead of just auditing access to files and shares, that type of thing, we can go in and audit when people change the SAM, audit the file system. I can audit the registry. That's a big one. Changing the registry is pretty tricky. Don't want people doing that.

Audit Registry

6:44-7:00

If you are going to Audit changing the registry, you also need to come down and click on Global Object Access Auditing and turn on Registry in here as well. We've got more with Privilege Use. Some system things we can audit, so lots of things in here that you can take a look at.

Audit Active Directory
7:01-7:14
Any time you see the letters DS like that, that's Directory Service, which means Active Directory. Now, I can audit Active Directory Replication or Active Directory Access, Active Directory Changes.
Those are my audit policies.
auditpol.exe Command
7:15-7:44
A couple of main things to take out is, if I'm setting up my Advanced Audit Configuration, and there are computers that are running Windows Server 2008, then I'm going to have to go through and still run auditpol.exe on those computers for them to be able to process the Advanced Audit Policy Configuration. If it's 2008 R2 or better, it should be able to handle Group Policy, no problem.
The last thing we're going to look at is that auditing feature that we can set up.
Set Up Auditing Example
7:45-8:37
If we want to Audit Object Access, we're looking at the second step where we go to the object itself and indicate what we're interested in auditing. Let's find a file or folder, and we'll just pick Test. This is actually done in the Advanced Securities, and you can see in here, there's an Auditing tab. I can Add an audit entry. First, I select who I suspect. It could be everyone, or it could be a particular user or group, and then I can go through and I say whether I'm looking for Success, Failure, or Both.
Finally, I specify what types of access I'm interested in. If Shad is getting in and writing, "The boss is funny looking with big ears," I'm going to Audit the Success of--at least, Modify--because modify is when we save or change data.
Adding a Condition
8:38-9:21
New with Windows Server 2012, not only does the interface look very different, but I also could set up a condition where I could say, well, I'm interested in whether the User or the Device is a member of a particular Group. They could be a Member of each group, and I could add five groups. It could be a Member of any of those groups, or I could say I'm only interested in auditing if they're Not a member of any of those groups.
Maybe I know there are three groups that should have access to this folder, and I'm simply interested in success of anyone who's not a member of those groups because that's a security issue, and that's how we set up auditing on the actual objects.
The results of my auditing are going to show up in the Security log.
Results of Auditing
9:22-10:03
We go into Event Viewer, Windows Logs, Security, and that's really the only thing that shows up in the Security Log, and you can see, we have a bunch of entries in here because there's some default auditing going on for our domain controllers, and those are the entries that are coming in.
Again, if you're looking for people getting logged on, you're going to have to look at all the domain controller logs, but you could go forward and set up a subscription so that you could centralize it on one computer that's monitoring the logs.
I hope you don't have to do a lot of auditing but, if you do, I'm glad that you know how to do it, and you can set it up so you can make sure that security is working properly in your environment.

8.4.3 Audit Policy Facts

In Windows, auditing records system events and other system changes. Auditing is enabled by configuring *audit policies*, either on a local system or through Group Policy. An audit policy is either enabled or disabled. When enabled, you choose to:

- Audit **Success** to identify who has gained access or who was able to exercise a right or privilege.
- Audit **Failure** to identify patterns of attempted access.

The following table describes the nine basic audit policies configurable through Group Policy.

Audit Category	Trigger Event(s)
Account logon	Account logon auditing tracks when a user account is used to authenticate to a computer. Account logon auditing generates an event on the system where the user account exists. • When a local user account is used, the local computer records the logon event. • When a domain user account is used, the domain controller records the logon event. In a multiple domain controller environment, you do not know which domain controller will authenticate a user. Event log subscriptions allow you to centralize the event log by collecting copies of specified events from multiple computers.
Account management	Account management auditing tracks changes to user accounts, including: • Create • Rename • Disable/enable • Delete • Change the password
Directory service access	Directory service access auditing tracks changes to Active Directory objects. Beginning with Windows Server 2008, Directory service access auditing capabilities have been integrated with Group Policy. The audit directory service access policy is divided into four subcategories: • Directory Service Access • Directory Service Changes • Directory Service Replication • Detailed Directory Service Replication

	When you enable Directory Service Access auditing, auditing for all four subcategories is enabled. To enable auditing for individual categories, use the **Auditpol /set /subcategory** command. When configuring directory service access auditing, enable auditing on the domain or OU, then identify the users and objects to audit. Simply enabling auditing using a GPO will be insufficient. To record the old and new values for changed objects, audit directory service changes. Auditing the directory service access subcategory creates a log entry when a change is been made, but does not log the actual values that were changed.
Logon	Logon auditing tracks logon or log off on the local system, or when a network connection is made to a system. For logon auditing, an audit event is recorded in the audit log of the local system, regardless of the type of user account used. For example, when a user logs on to a computer using a domain account, a logon event is recorded on the local workstation, while an account logon event is recorded on the domain controller.
Object access	Object access auditing tracks access to files, folders, or printers. You can also audit actions taken by a certificate authority, access to specific registry settings, or access to specific IIS metabase settings. For file auditing to occur, the files must be on NTFS partitions. In addition to enabling auditing in the audit policy, you must configure auditing on the specific objects you want to track.
Policy change	Policy change auditing tracks changes to user rights, trust relationships, IPsec and Kerberos policies, or audit policies.
Privilege use	Privilege use auditing tracks the following actions: • A user exercises a user right. • An administrator takes ownership of an object.
Process tracking	Process tracking auditing records actions taken by applications. Process tracking auditing is used mainly for program debugging and tracking.
System	System events auditing tracks system shutdown, restart, or the starting of system services. It also tracks events that affect security or the security log.

Be aware of the following when configuring auditing:

- Auditing requires system resources.
- You view audit entries in the Event Viewer Security log.
- In Windows Server 2012, you can set up conditional auditing.

- With both Directory Service Access and Object Access auditing, configuring auditing requires two steps:
    1. Enable auditing in the local security policy or Group Policy.
    2. Configure auditing on the specific objects.

    View the System Access Control List (SACL) of the Active Directory object or the NTFS file or folder to identify the users, groups, or actions to track.

- In addition to tracking the necessary events, make sure your logs are properly configured to save all of the necessary information.
    - Use the Event Log policies in Group Policy to configure the Security log size and retention method.
    - To preserve all logged actions, configure logs to not overwrite events. When logs are not configured to clear automatically, you must periodically save and clear the logs to make room for additional events.
    - Enable the **Audit: Shut down system immediately if unable to log security audits** security option to prevent the system from being used if the log is full (this setting is also referred to as CrashOnAuditFail).

Beginning with Windows Server 2008 R2, advanced auditing capabilities were integrated with Group Policy. Advanced auditing offers 53 setting that allows you to eliminate unwanted data and specifically target data important for system management and security. Advanced auditing settings can be used in place of the nine basic auditing settings. If you use Advanced Audit Policy Configuration settings, enable the **Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings** policy under **Local Policies\Security Options**. This will prevent conflicts between similar settings by forcing basic security auditing to be ignored.

8.5 User Rights Assignment

As you study this section, answer the following questions:

- What is the difference between *permissions* and *rights*?
- What tool do you use to configure rights policy settings?
- By default, what groups have the **Allow Log On Locally** right?

After finishing this section, you should be able to complete the following task:

- Configure user rights to secure network resources.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  - Manage Security Policies
    - Configure User Rights

This section covers the following 70-410 exam objectives:

- 502 Create and manage Active Directory users and computers.
  This objective may include but is not limited to:
  - Configure user rights
- 602 Configure security policies.
  This objective may include but is not limited to:
  - Configure User Rights Assignment

8.5.1 User Rights

User Rights
0:00-0:02
I want to take a few minutes to talk about user rights.
Permissions
0:03-0:14
A lot of people get confused between rights and permissions. Permissions are my ability to use objects such as files, folders, printers, any type of an object, it could even be an OU.
Rights
0:15-0:27
Rights are my ability to perform some type of action within the computer. I have the right to shut down the computer. I have the right to change the system time. I have the right to back up files. I have the right to restore files.
Rights and Group Policy
0:28-0:46
Most network administrators are used to working with groups, -- either local groups if it's a member server, or domain local groups if it's on a domain controller, because these groups have already been granted the appropriate rights in Group Policy. Those groups are getting those rights from the user rights section of Group Policy.
Summary
0:47-0:56
Just as an overall rule, permissions are for objects, rights, or for activities on the computer, and rights come from Group Policy in the user rights node which we'll take a look at in the demo.

8.5.2 Managing User Rights

Managing User Rights
0:00-0:24
In this video, we're going to take a look at managing user rights. You may already know that there are groups that exist in Windows that have particular rights assigned to them. For example, the Backup Operators group has the right to back up files and folders. On a local workstation, the Backup Operators group in the SAM has rights to that particular member's server or client. On a domain controller, backup operators will have rights to backup files on all of the domain controllers.

Group Policy
0:25-0:36
These rights actually come from Group Policy, so we're taking a look at a member server, and we're going to look at the Local Group Policy, but you can assign rights using any level of Group Policy that would be applicable.

Example's of Rights
0:37-0:53
We'll get into the Local Group Policy by typing gpedit.msc. I'm going to right click it and Run as Administrator. My user rights are always in Windows Settings, Security Settings, Local Policy, User Rights Assignment.

Rights
0:54-1:30
Rights are my ability to do things to the computer. For example, I could Access this computer from the network. I could logon locally or logon through Remote Desktop Services. Here's the right to Backup files and directory, and we can see that Backup Operators has been assigned that right to logon locally and to Backup files and directories. Backing up files and directories is not the same as restoring them, so there's a corresponding right down below to Restore files and directories which has also been granted to Backup Operators.
Another right is the ability to Shut down the system.

Rebooting a Server
1:31-2:38
We don't want just anybody to be able to reboot the computer. Anytime you have to reboot a server you should be if not a little bit afraid then a lot afraid. The worst I've ever seen it, it took five days for a server to reboot. A lot of times, viruses and spyware need a reboot to take effect, so you want to make sure you always have a good backup, and you're prepared, and it's after hours, so that you know the server will be coming back in time so that users don't have any problems.
You don't need to memorize all of the local user rights. Simply be familiar with the types of things they do. They allow me to perform functions on the computer like access it, back it up, restore it, remove it from a docking station. We can even Deny logon locally or Deny logon through Remote Desktop Services. Suppose there was a particular user account that is in a group that's allowed to remote desktop into the computer, but just that one account shouldn't be able to do that. We need to deny that particular account the ability to get in through remote desktop, but we can't remove them from the group. I can actually proactively assign them the Deny logon through Remote Desktop Services right, and even though they're a member of a group that has that right normally, they would be denied.

Summary
2:39-2:49
User rights allow users to perform functions on the computer. They're given out with Group Policy inside the local policies, and we can adjust them at any level of Group Policy that makes sense.

8.5.3 User Rights Facts

Permissions are the ability to use objects, such as files, folders, and printers. *Rights* are the ability to perform actions on a computer, such as log on, shut down, back up, and restore. For example, a user logging on locally must have the **Allow Log On Locally** right.

Rights are applied locally to individual computers. You can view the User Rights Assignment policy settings in Computer Configuration\Security Settings\Local Policies in Group Policy. When assigning rights, keep the following in mind:

- Rights are part of the security policy for the computer.
- Rights can be assigned using local or domain policies. Use the Local Group Policy Editor (for local policies) or Group Policy Management (for domain policies) to configure user right policy settings.
- If a right is assigned in a domain GPO, the right affects the local security settings of the computer accounts to which the GPO is enforced.
- The Default Domain Controllers Policy GPO in Windows Server 2012 assigns the **Allow Log On Locally** right on domain controllers to the following groups by default:
  - Account Operators
  - Administrators
  - Backup Operators
  - Print Operators
  - Server Operators

  For workstations and member servers, the Allow Log On Locally right is assigned to the following groups by default:

  - Administrators
  - Backup Operators
  - Power Users
  - Users
  - Guest
- You can explicitly deny a right to users or groups. For example, you could deny the Print Operators group the right to log on locally.

There are many User Rights Assignment policies that can be used to manage what users are allowed and not allowed to do on the system where the policies are applied, including:

- Access this computer from the network
- Add workstations to domain
- Allow log on locally
- Allow log on through Remote Desktop Services
- Back up files and directories
- Change system time
- Force shutdown from a remote system
- Load and unload device drivers
- Manage auditing and security log
- Perform volume maintenance tasks
- Profile system performance
- Restore files and directories
- Shut down the system
- Take ownership of files or other objects

8.6 Security Options

As you study this section, answer the following questions:

- What is the best practice for managing Administrator and Guest accounts?
- How would you prevent network users from accessing resources on an optical disc?
- How can you standardize security settings across multiple computers on a network?
- What is the role of the User Account Control (UAC)?
- What is the difference between *prompt for consent* and *prompt for credentials*? When would one or the other be displayed?

After finishing this section, you should be able to complete the following tasks:

- Configure security options.
- Enforce User Account Control.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Manage Security Policies
  - Configure Security Options
  - Enforce User Account Control

This section covers the following 70-410 exam objective:

- 602 Configure security policies.
  This objective may include but is not limited to:
  Configure Security Options settings
  Configure Security templates
  Configure Local Users and Groups
  Configure User Account Control (UAC)

8.6.1 Security Options

Security Options
0:00-0:17
Security options are a subset of Group Policy that contain all the security options. In there are my user rights, my settings for the UAC, anything that might have to do with security, Windows Firewall--all of that stuff is in that subcategory.

Security Templates
0:18-0:57
If you have settings that you want to be standardized, you can actually export that whole node in Group Policy into a file. We call those Security Templates. Then what you would do with that Security Template is import it into another Group Policy so that you can roll it out to more machines.
If I had a security standard for my company that says, "Okay, these are all the things that need to be turned on. These are things that need to be turned off. This is what the Firewall would look like." I would go into one computer, modify all the Group Policy Options under Security Settings, export that as a Security Template, and then go into the appropriate policy, maybe the Default Domain Policy, and import the template.

Configuring Security Options
0:00-0:44
In this video, we're going to take a look at configuring security options. Security options can be configured at any level of Group Policy, whether it's Local Policy or it's policies up in Active Directory. We're going to take a look inside Group Policy Management Console. I'm going to get it out of the Tools menu. We can certainly go the Start menu. Since we're just interested in settings, I'm going to edit the Default Domain Policy. We need to go ahead and expand Policies, Windows Settings, Security Settings, Local Policies, and then we want to go into Security Options.

User Rights and Security Options
0:45-1:17
User rights are my right to do something on the computer. Security Options are exactly that, we're controlling security options essentially.
This isn't really governing necessarily what the users can do, it governs what the rules are for security on this particular computer. There's a lot of settings in here-you don't need to memorize all of them. I want to point out a few that can be pretty important. You have the Administrator Account Status.

Accounts: Administrator Account Status
1:18-2:21
By default, the Administrator Account is disabled on client workstations, and that's the way it should be.
If you have older workstations like XP, you can come in here and Enable the account status, and then I can go through and set it to Disabled. You can see it determines whether it's Enabled or Disabled. Actually what I would be doing is disabling the administrator account, but if you did need to enable it because it's disabled, then you can enable it as well. It's disabled because it's considered a security risk. Anybody who's ever installed any kind of Microsoft operating system knows there's an account named administrator. This is the most powerful account on the computer.
They've got 50% of the information that they need to hack the most powerful account on the computer. If you don't want to disable the Administrator account or leave it disabled, at the very least you might want to rename it so that if somebody is going to try to hack that account, it's not named Administrator anymore--it's named something else.

Accounts: Rename Administrator Account
2:22-2:41
I can also rename the guest account--enable or disable the guest account as well.
Down in devices, I can prevent users from installing printer devices.

Devices Security
2:42-3:04
I can restrict CD-ROMs and floppy devices to just the user that's logged on if you want to do that, so that people can't come in across the server and use those. We also can say whether or not they're allowed to format and eject removable media.
Big ones that you would see in a domain would be here in Interactive Logon.

Interactive Logon: Display User Information When the Session is Locked
3:05-3:56
Display user information when the session is locked. If somebody is coming through, walking through your company even though the display is locked, they can see the username. Again, if I get usernames, I've got 50% of the information that I need to hack accounts.
You would probably want to turn off Display user information when the session is locked; I would also turn on Do not display last username. The user logs out, goes home for the day, somebody comes in there in the building for whatever reason; they can't just hit CTRL + ALT + DEL and there's the name of the last user that logged on, they've got to start from scratch.
The user will cry and say it's annoying, but it's definitely helpful for security.

Interactive Logon: Do Not Require CTRL + ALT+ DEL
3:57-4:29
Do not turn on "Do not require CTRL + ALT + DEL", that is a security issue. That's been called for many years the SAS, or Security Attention Sequence, and what it does is when you hit CTRL + ALT + DEL and that logon screen comes up, it stops any software that would be running in the background. That's to prevent Trojans so that there can't be any Trojans running while they type in their username and password, and the Trojan would capture that.
If you turn that off, then you're essentially disabling that protection.

Interactive Logon: Message Text for Users Attempting to Log On

4:30-5:18

Another one we see very often would be Message text and Message title, for users attempting to log on. So if you've ever logged on to the screen and a box comes up and says "warning, this is just for authorized users of the northsim.com company, if you're not a northsim.com company, don't log on", that would be the message text.

The story I've heard that very early in the days of hacking, there was a hacker that was taken to court and basically their excuse in court was, well, I saw the Welcome to Windows box, I thought it was welcoming me and I just needed to figure out my username and password. So you don't have protection from that type of argument in court unless you have this message text. You see it in a lot of companies, and then in other companies it's not so big.

Interactive Logon: Prompt User to Change Password Before Expiration

5:19-5:27

I'd also turn on "Prompt the user to change the password before expiration". The odds of them remembering to change their password without being prompted is not very good.

Interactive Logon: Require Smart Card

5:28-6:13

We often will see "Require smart card". Smart card is, sometimes they look like cards, sometimes they don't look like cards, but it's some type of physical token that has to be inserted into the computer, and it has a certificate on it that uniquely identifies that user.

The idea being that now I've got two factor authentication, because not only do they need a username and password, they also need to have this card as well. If they write down their username and password somewhere, still somebody would have to steal their smart card in order to get in. You see this in various secure environments. The great thing about it is it's much more secure. The bad thing about it is you go to buy hardware and smart cards for everybody, so it can be expensive to implement.

Got a whole bunch of Network access settings in here.

Network Access Security

6:14-6:25

These I don't see used that frequently. See, there's very few of these that are enabled or disabled by default.

Network Security: Force Logoff When Logon Hours Expire

6:26-7:01

Another great one is "Force logoff when logon hours expire". This is disabled by default, the idea being if I'm only allowed to log in between nine and five but I've got to work late, it's not going to boot me off at five, it just means if the computer does happen to reboot, I won't be able to log back in.

If you are really trying to enforce logon hours, then you would want to come in and turn this on, and that makes sure that not only can they not log on after hours, but they can't remain logged on and stay there until midnight when they try to hack everything.

Shutdown: Allow System to be Shut Down Without Having to Log On

7:02-7:16

This will set up whether they're allowed to shut their system down without having to log on. On workstations that usually can do that, on the log on screens on servers, we usually can't. If you want that to be available on servers, then you can certainly go ahead and turn that on.

User Account Controls Settings

7:17-7:20

Then we've got the User Account Control Settings down here.

Summary

7:21-7:57

Again, you don't have to memorize everything in here--a lot of great settings. Certainly something to look through, and you're always trying to balance off having the computer be very secure but not having it be so secure that it's not functional for the users. Pick and choose wisely from this list. These are the security options that you can use to help control the security of the workstations and servers and set up a baseline security policy so that your network is going to function as well as it possibly can.

8.6.4 Security Options Facts

Security options are a subset of Group Policy that governs the rules for security on the computer.

The following table identifies important Group Policy security options grouped by the setting category.

Setting Category	Description
Accounts	Be aware of the following regarding security options in the Accounts category: • The **Administrator account status** policy specifies if the administrator account is enabled or disabled. By default, the Administrator account is disabled on client workstations. If the computer starts in Safe Mode, the Administrator account is enabled. • It is best practice to: Disable the Administrator account and Guest account. If you cannot disable the Administrator and Guest accounts, rename them using the **Rename administrator account** and the **Rename guests account** policies. Enable the **Limit local account use of blank passwords to console logon only** policy.
Devices	Consider using the following policies to secure devices based on the security needs of the organization. • **Prevent users from installing printer devices** protects the system from the possible introduction of incompatible drivers or drivers infected with malware. • **Restrict CD-ROM drive access to locally logged-on user only** prevents network users from accessing resources on a CD-ROM. • **Allowed to format and eject removable media** protects removable media from users when disabled. • **Unsigned driver installation behavior** specifies what happens when an attempt is made to install an unsigned driver.
Interactive logon	Best practices for the following Interactive logon settings are: • Disable **Display user information when the session is locked**. • Enable **Do not display last user name**. • Disable **Do not require CTRL+ALT+DEL.** The key sequence has been referred to as the security attention sequence. It stops any software running in the background to prevent Trojans from capturing the user name and password.

	• Use the **Message text for users attempting to log on** to provide a disclaimer for that the computer is only for use of company employees. • Enable **Prompt user to change password before expiration** as a reminder to the users to change their passwords. • Enable **Require smart card** is an additional security measure.
Network security	Consider using the following policies to secure network access based on the security needs of the organization. • Enable **Force log off when logon hours expire** to enforce logon hours. • Disable **Allow system to be shut down without having to log on** to require a user to log in before shutting down a server.

To standardize settings, you can configure the security options and then export the Security node to be used as a Security Template. You can then import the Security Template into Group Policy on another computer. GPO templates are settings that you copy or import into new GPOs to enforce common settings.

8.6.5 User Account Control

User Account Control
0:00-0:21
Now we're going to talk about User Account Control. This feature came in with Windows Server 2008 and Vista on the client side. Here we've got a basic client, I've got my domain controller, I've got a folder maybe on the domain controller or on a file server, it doesn't matter, and then on the Access Controller List for that folder, the only people who have rights are domains Admins, and they have full control.

Administrator and Standard User Account
0:22-1:30
Since I've started teaching and working in the industry, which is almost 18 years now, Microsoft has always been telling Admins, "What you need to do is make yourself an Administrator Account and a Standard User Account." Go in and use your Standard User Account when you're doing things as a User. When you need to do something as an Administrator, right click the application, do a Run As, now we have Run as administrator, and provide your Administrator credentials. The purpose of this is that we don't want to go through surfing the Internet or doing something like that. A virus hits, or a spyware hits and now I'm logged in as an Administrator, and that virus or spyware has full administrative rights to the Domain to install itself in my computer and maybe in other computers in my environment. Unfortunately, not a lot of people followed this. I think I've worked at maybe one company that actually had that as a policy. Most of the places I've worked have just had one account. It's a member of Domain Admins or Enterprise Admins, and that's what I use both for my day-to-day-- sending e-mails, surfing the Internet, researching problems-- and for doing administrative work.

UAC
1:31-1:37
Starting with Vista in 2008, Microsoft implemented the UAC, which pretty much forced you to do this.

Example Without UAC
1:38-1:54
What happens in a normal situation is this: I sit down at my client, and I log in with my administrator account. Let's say my Shad account is a member of Domain Admins, my computer is going to send out to the domain controller and say, I have this geek, he says his name is Shad, his password is password, is that a good match?

Access Token
1:55-2:00
If it is, I'm authenticated, and the domain controller sends me down an access token.

Security ID (SID)
2:01-2:46
You can't see the access token, but what's on it is this: my user name, my password, and then the SID for my User Account. Because I put in the name Shad, the computer doesn't know me as Shad, they know me as my security ID, a big long number that identifies me uniquely, also the SIDs of any groups to which I belong.
All that's kept on my access token. When I go to access an object like this folder over here, on the folder is the access control list, it's usually the security tab, my computer takes that token and compares it to the list on the object. If there's a match between my SID and any group SIDs, I'm going to get into that object. If there's not a match, I don't get in. There's an implicit deny. This is what happens in normal circumstances before the UAC.

Example With UAC
2:47-2:50
Let's take a look at how this changes when we put the UAC into the picture.

Two Types of Access Tokens

2:51-3:55

Here I have the same situation, I log on with my User Account, Shad, it's sent up to the domain controller, I get authenticated, but now what comes back are two access tokens. One access token is a standard access token, so it's got my username, my password, my SID, my group SIDs except for any administrative group. So administrators, Domain Admins, Enterprise Admins-- anything that's an administrative SID is on the other access token. So all these SIDs for the groups--I'm just going to write Admins, but it could be any of those: administrators, Domain Admins--it's all on that other token.

Essentially what happens is this: if I go to access something and it's got User Privileges, I just get in. When I go to access a resource that only has rights for administrators, so here only Domain Admins have full control, the computer's got to hop over from my regular token and send that (the administrative token) and that's what triggers the UAC. So it's going to come up with a box and say, "Hey, you've got to be an administrator to do this, do you want to continue? OK or Cancel." That's essentially what the UAC is.

Benefits of the UAC

3:56-4:28

The great thing about the UAC is that whenever I try to do something that requires me to be an administrator, it has to hop over to that token, I get the prompt, and now I'm aware of what's going on. I go out to a website, it wants to install a virus. Now, at the very least, it's going to trigger that consent box and say, "Hey, something is trying to do this, it requires administrative rights. OK or Cancel." Now I'm aware an administrative change is being made on the computer; I didn't authorize it, it's nothing I'm intending to do, and I can say cancel.

Secure Desktop

4:29-5:24

The other thing you want to keep in mind with the UAC is when that dialogue box comes up, it dims the screen. There was a lot of backlash against the UAC in Vista. Part of it was, people weren't used to all these dialogue boxes popping up and saying, "Hey, you want to say OK or Cancel," every time they tried to do something, it requires them to be an administrator. The other thing that didn't go over well with Vista was this dimming of the screen. That's actually called the Secure Desktop, and it's a really good feature that improves the security that's added by the UAC. When the screen dims, and you're on the Secure Desktop, it stops all other processes running in the background, except the UAC. What that means is the virus makers or the spyware makers can't develop spyware or a virus that would click OK for you. That secure desktop stops everything except the UAC-- if somebody's clicking OK, it's definitely the end user.

Configuring User Account Control
0:00-0:19
In this video, we're going to take a look at configuring User Account Control.
I need to get into Group Policy Management Console. I'm going to go in through the Tools menu. Since we just want to look at the policy, I'm just going to edit the Default Domain Policy. We need to expand Policies, Windows Settings, Security Settings, Local Policies, and then you want to click on Security Options.

Security Options
0:20-0:41
I always remember it's at the very end of this category.
We can see that there's a number of settings down here that have to do with User Account Control.

Admin Approval Mode for the Built-In Administrator Account
0:42-1:00
This particular one governs what will happen when the Built-in Administrator account runs some type of an application. It will make sure that even if it's the Built-in Administrator, it will trigger Admin Approval Mode, and then whatever's set up for Admin Approval Mode will be inflicted on that account.

Behavior of the Elevation Prompt for Administrators
1:01-1:31
The two that we're most interested in are the ones that start with "Behavior of the elevation prompt for" and there's two of them. One is administrators in Admin Approval Mode. This goes with this. The top one says, "the Built-In Administrator account is going to be subject to Admin Approval Mode."
And then this setting says what's going to happen in Admin Approval Mode. There's a few choices as to whether we're going to see the Secure Desktop or not.

Prompt for Consent and Prompt for Credentials
1:32-2:34
We have Prompt for credentials and Prompt for consent. Prompt for consent is the one that you're probably used to seeing, which is OK or Cancel. Prompt for credentials is more secure. It means if I go to do anything as administrator, not only am I going to get a prompt, but I'm going to have to re-enter my administrative credentials. That, I would agree, can probably be a little bit annoying. I would only be using that in a really secure environment. Those are the ones that you want to be aware of.
Prompt for consent, it just lets me say OK and keep going, or Prompt for credentials, which forces me to enter in my username and password. If you had a scenario where company policy says that everybody should have to enter in their username and password before installing software or doing anything that requires administrative rights, then that's code for, "even administrators should be prompted for credentials." This policy governs administrators.

Behavior of the Elevation Prompt for Standard Users
2:35-3:25
This policy governs standard users. Standard users are not administrators. They're going to be prompted no matter what, because if they go to do something as an administrator, they don't have any rights.
You've got a couple of choices. By default, it says "Prompt them for credentials," and I can choose a secure desktop or not. If this were my environment, I'd probably be doing "Automatically deny elevation requests." The reason being, if you prompt them for credentials they think, "If I could just browbeat Shad into telling me the right username and password, I can have this." Automatically deny says "no, I'm sorry, this is for administrators only," something along those lines-- that you have to log the ticket or contact their system administrator. It just sends more of a message of "this is not something you can do. This is something that belongs to administrators. Stay out of it."
We also can have it Detect application installs and prompt for elevation.

Detect Application Installations and Prompt for Elevation
3:26-3:38
If you disable that, the applications will be able to be installed without them providing administrative rights. Not necessarily recommended, but certainly possible to do.

Only Elevate Executables that are Signed and Validated
3:39-3:47
If I want to make sure that all of the executables that are being installed have been signed and validated, I can turn on this setting here.

Run All Administrators in Admin Approval Mode

3:48-3:56

Here, I can say I want to Run all administrators in Admin Approval Mode; whereas the top one just talks about the Built-in Administrator Accounts, this talks about all administrator accounts.

Virtualize File and Registry Write Failures to Per-User Locations

3:57-4:27

The last one is one that you may also have to interact with. What it does is this: When applications try to write to protected areas like Windows System 32 or program files, historically, they used to be allowed to do that. The problem came in, if you accidentally get a virus or spyware on the machine, it may actually install itself into areas that we don't want files being modified. I believe this came in with Server 2008 in Vista.

Virtual Locations

4:28-5:25

What happens is there are virtual locations in the user profile and we can't see all of them. You can see one of them here-- Application Data, Local, Roaming. Instead of writing into Program Files, it redirects them to areas like this Application Data in the Profile. The application thinks it's saving into a sensitive area like Windows or System 32. It's actually being saved into an area of the Profile so the operating system doesn't become corrupt. If you have something that doesn't work with that, you can use that very last setting to disable that. Anything that you're turning off in the UAC, you're making the computer less secure. Those are some of the settings you can set up if you need to configure the behavior of User Account Control.

8.6.7 User Account Control Facts

User Account Control (UAC) helps minimize the dangers of unwanted actions or unintended software installations. UAC insures that actions which affect the system configuration are approved by users with the necessary rights to perform those tasks. To understand how UAC works, be aware of the following accounts:

Account Type	Description
Standard user	A *standard user* account has the least amount of user rights and privileges required to perform most basic tasks.
Administrator	An *administrator* account can perform any action on the computer. For example, administrators can turn off firewalls, configure security policy, and install new drivers and other software for the entire computer. • Administrators are members of the local Administrators group. • Each local computer has a built-in administrator account that exists by default. • During a new installation, the first user account you create is an administrator account. Subsequent user accounts are standard users. The built-in Administrator account is then disabled. • For upgrades, the built-in administrator account is enabled if it is the only user account with administrative privileges, otherwise it is disabled. • If the system has at least one administrator account, the built-in administrator account cannot be used to log on to Safe Mode. Safe Mode logon using the built-in administrator account is never allowed for computers that are members of a domain.

When a user logs on to the system, an *access token* is generated for the user. The access token controls the type of actions that the user can perform on the system. The default behavior of UAC is as follows:

- The access token identifies the user account as either a standard user or an administrator. Certain actions can only be performed by a user with an administrator access token.
- When a standard user logs on, a standard user access token is generated. When an administrator logs on, two access tokens are generated: one standard user token and one administrator token.
- The standard user token is used to attempt to perform all tasks for both standard users and administrators.
- If standard user rights are not sufficient to perform the task, the system requests *privilege elevation*:
  The standard user is prompted to provide administrator user credentials (username and password). This process is referred to as *prompt for credentials*.
  The administrator user is asked whether the administrative token should be used to perform the task. Because the administrator has already logged on with the username and password, this is a simple Continue or Cancel question. This process is referred to as *prompt for consent*.

- Using a standard user access token and prompting for consent for administrators is referred to as Admin Approval Mode. This feature of UAC helps protect the system by running all processes using the least administrative privileges necessary.
- Prompting for credentials or consent activates the Secure Desktop. With the Secure Desktop, the desktop and all active applications are darkened, and the prompt appears over the shaded desktop. You must respond to the prompt before you can continue with the requested operation or return to the desktop. The Secure Desktop prompt will be displayed for 150 seconds, after which the request for privilege elevation is automatically denied.
- Although Admin Approval Mode provides some degree of protection, it only prompts for consent without asking for a password. If the user is logged on as a standard user, the Administrator username and password is required.

The process described above is the default behavior of UAC. You can customize many aspects of how UAC works.

Use the UAC settings in the Control Panel to configure the sensitivity of UAC. You can adjust the UAC configuration to different levels of notifications to reduce the constant or unnecessary UAC prompts. Notification level settings include the following:

Setting	Details
Always notify	• A UAC prompt and the Secure Desktop is displayed for 150 seconds. • The user cannot perform any other actions until a response to the prompt is entered. If there is no response after 150 seconds the request is automatically denied. This is the most secure and recommended configuration.
Notify me only when programs try to make changes to my computer	• Prompts only when programs try to make changes to the computer or Windows settings. • A UAC prompt and the Secure Desktop is displayed for 150 seconds. • The user cannot perform any other actions until a response to the prompt is entered. If there is no response after 150 seconds the request is automatically denied.
Notify me only when programs try to make changes to my computer (do not dim the desktop)	• Prompts only when a program is trying to make changes to your computer or when a program that is *not* included with Windows attempts to modify Windows settings. • The Secure Desktop is not displayed.

Never notify	• If logged on as an administrator, all actions are executed without UAC prompts or the Secure Desktop. • If logged on as a standard user, all actions requiring privilege elevation are automatically denied. Turning UAC off requires a system reboot.

The following table describes the equivalent Group Policy settings for each notification level.

Setting	Group Policy Equivalent
Always notify	Use the following Group Policies for the equivalent of **Always Notify**: • The **Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting is set to **Prompt for consent on the secure desktop**. • The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is enabled.
Notify me only when programs try to make changes to my computer	Use the following Group Policies for the equivalent of **Notify me only when programs try to make changes to my computer**: • The **Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting is set to **Prompt for consent for non-Windows binaries**. • The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is enabled.
Notify me only when programs try to make changes to my computer (do not dim the desktop)	Use the following Group Policies for the equivalent of **Notify me only when programs try to make changes to my computer (do not dim the desktop):** • **The Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting is set to **Prompt for consent for non-Windows binaries**.

	• The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is disabled. • The **Behavior of the elevation prompt for standard users** policy setting is set to **Prompt for credentials**.
Never notify	Use the following Group Policies for the equivalent of **Never notify**: • The **Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting is set to **Elevate without prompting**. The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is disabled. • The **User Account Control: Run all administrators in Admin Approval Mode** policy setting is disabled. • UAC is disabled. If you use Group Policies to turn off UAC, reboot the system for changes to take effect.

8.7 Restricted Groups

As you study this section, answer the following questions:

- When is it a good idea to create a Restricted Group Policy?
- What are the two methods to define a Restricted Group? Which is the preferred method?
- Why is it a good strategy to test before activating a GPO using a Restricted Group Policy?

After finishing this section, you should be able to complete the following task:

- Configure restricted groups.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Manage Security Policies
  - Configure Restricted Groups

This section covers the following 70-410 exam objective:

- 503 Create and manage Active Directory groups and organizational units (OUs).
  This objective may include but is not limited to:
  Manage group membership using Group Policy

8.7.1 Restricted Groups

Restricted Groups
0:00-0:16
We're going to talk about a specific section of Group Policy called Restricted Groups. What the Restricted Groups node allows us to do is to control membership of groups on either the client or the member server using a Group Policy. There's really two ways to do this.

Add a Local Group and Then Define the Membership
0:17-0:40
One way is to add in a local group into Restricted Groups and then define the membership of that local group. For example, I can add in administrators, and then define who should be a member of administrators. When that policy is applied to the client or the member server, anybody that's not in my policy will be removed. Anybody that's in my policy that's not in administrators will be added.

System Groups
0:41-0:55
That's a little risky if you're going to use it to control the default groups like users and administrators, because you have to make sure that you don't accidentally remove a system group that the computer needs in those local groups in order to run.

Add a Group on the Domain Which is Added to a Local Group
0:56-1:44
What's a little bit easier to do is to use the other side of Restricted Groups and add in a group up on the domain that would get added to a local group.
For example, if I wanted to make a group called Desktop Admins and have that be added to the administrators group on all the desktop machines, I can apply my policy to the desktop I'll use, add in the Desktop Admins group, and specify that it will become a member of the local administrators group.
A little bit less risky, because all that policy is going to do is add my group to that administrators group on the client or the member server. It's not going to take anybody out of it. If you do have a very secure environment where you need to make sure that membership of local groups is limited to particular accounts, then certainly Restricted Groups will do that.

Managing Groups with Group Policy
0:00-0:04
In this video, we're going to take a look at managing groups using Group Policy.
Group Policy Management Console
0:05-0:39
I want you to go into Group Policy Management Console. I'm going to go in through the Tools menu, but I can also use the Start menu as well.
Since we just want to look at the Settings, I'm going to Edit Default Domain Policy. We need to expand Policies, Windows Settings, Security Settings, and then the node we're interested in is Restricted Groups.
Restricted Groups Node
0:40-0:48
We can use Restricted Groups in one of two ways. Let's take a look at that. I'm going to right click and Add a Group.
Control Group membership on the Local Workstation
0:49-1:08
One way is to control group membership on the local work station. Maybe I want to control who has rights on any of the work stations in my domain. I can say, okay, I'm going to control the membership of the Administrator's group. Then it says, okay.
Define Who Should Be a Member of This Group
1:09-1:25
Define who should be a member of this group. By default, the group should contain no members, or if I need to add this group to another group, I can do that down here.
If it's a local group like Administrators, I really can't add it to another group. I'm kind of stuck up here.
Default Members
1:26-2:05
This is a cause for concern because, if you look at a member server or a Windows 8 workstation, you're going to see that there are default members of the Administrator group. For example, every machine that joins the domain, it adds the Domain Admin's group into the Local Administrator's group. If I don't re-create that membership in here, it would pull out all the members of that group, and now, we have no way to administer that workstation.
If you're going to specify local groups, you need to be very careful and see who should be in that group and replicate it up in Group Policy, or you could actually run into a problem.
This Group is a Member Of
2:06-3:17
A better way of doing it would be to use This group is a member of. In that case, I can't use a local group. I've got to use a group up in Active Directory. I could come in and say, well, maybe I have a situation where I want to add a group to the Local Administrator's group so that I can give out administrative rights on the workstations that will be affected by this policy. I add a group and let's just say it's the Backup Operator's group, could be a group I create. It doesn't matter which group it is.
The domain group Backup Operators-- I'm going to go through and add administrators--and every machine that's affected by this policy is going to have the Backup Operators group added to the local administrator's group.
Make sure that these policies, if I'm leaving this area blank, are not going to be applied to the domain controllers. I probably would not be doing this in the Default Domain Policy. I would be doing this in a policy that affects a bunch of member servers of a bunch of clients, but when that policy hit, it would take Backup Operators and add it to Local Administrator's group.
Summary
3:18-3:38
The way we control group membership using Group Policy is to use the restricted groups node in Group Policy. Again, we want to make sure that we accommodate for whoever should be a member of that group by default, and we want to make sure that we apply the policy on a scope such that we don't accidentally remove anybody who really should be a member of that group.

8.7.3 Restricted Group Facts

The Restricted Groups Policy is a powerful tool that can be used to control membership for groups that require high security. One potential use for Restricted Group Policy settings is managing the membership of local groups on domain member servers and workstations.

Using the Group Policy Management console, a restricted group can be defined in two ways:

- **Members of this group** identifies individual members of a restricted group. All users listed become members of the specified group on systems where the policy is in effect.
  Any user who is currently a member of the group but whose name is not on the list is removed from group membership.
  Any user on the list who is not currently a member of the restricted group automatically becomes a member.
- **This group is a member of** defines one or more groups the restricted group becomes a member of.

Use this option to define membership in a local group by adding a restricted group. The restricted group to be added to the local group must be a group defined in Active Directory.

Using **This group is a member of** is the preferred method for defining membership in a restricted group.

When using the Restricted Groups Policy, keep in mind:

- Once an administrator has designated group membership with a Restricted Groups Policy, no one can add or remove members. A user can use other tools to change the group membership, but a refresh of the group policy settings will overwrite any changes made.
- The Restricted Groups Policy does not change group membership in other groups.
- When using the Restricted Group Policy to control membership in default local groups, carefully identify all system groups that the computer, applications, and legacy applications need to run. The implications of leaving out a critical user or group can be severe.
- If you link a GPO with Restricted Groups Policy settings to a domain, the setting will be inherited by all computers in the domain, including domain controllers and Active Directory security groups.
- Testing is recommended before activating a GPO using Restricted Groups Policy in a production environment.

8.8 Windows Firewall Rules

As you study this section, answer the following questions:

- When should you use the Advanced Firewall instead of the Basic Firewall?
- When would you configure a custom exception?
- What does the exception scope do?
- By default, which type of traffic is allowed through the firewall?
- What is the benefit of using connection security rules?

After finishing this section, you should be able to complete the following tasks:

- Use the Basic Firewall to allow traffic based on port, protocol, or application.
- Use the Windows Firewall with Advanced Security to manage custom firewall rules.
- Use Group Policy to enforce firewall rules.

This section covers the following 70-410 exam objectives:

- 604 Configure Windows Firewall.
  This objective may include but is not limited to:
    - Configure rules for multiple profiles using Group Policy
    - Configure connection security rules
    - Configure Windows Firewall to allow or deny applications, scopes, ports, and users
    - Configure authenticated firewall exceptions
    - Import and export settings
- 203 Configure servers for remote management.
  This objective may include but is not limited to:
    - Configure Windows Firewall

8.8.1 Inbound and Outbound Rules

Inbound and Outbound Rules
0:00-0:11
In this video, we're going to talk a little bit about the Windows Firewall. Microsoft added in a firewall as far back as XP. We started moving into Vista, and the later operating systems, they really beefed it up.

Basic Firewall
0:12-0:33
When you're working with the basic firewall, there's not that much you can do. You can turn it off or on, you can create exceptions for applications, but those exceptions will go in both directions. Then you can specify which network profile the exception will be attached to. Private networks, public networks, or domain, otherwise known as work, if you're in a domain environment.

The Windows Firewall with Advanced Security
0:34-1:24
The Windows Firewall with Advanced Security gives you a lot more flexibility, so you can go in there and have different rules for inbound traffic versus outbound traffic. We can specify the rule to affect a service, a port, a particular range of IP addresses, certain types of traffic, different types of ICMP packets. I can really get discreet about exactly what type of traffic the rule applies to and then set up what the computer will do. Is it going to allow this traffic, block the traffic, or allow the traffic only if it's secured using IPsec?
Of course, I can specify which network profile it will be attached to. The Windows Firewall with Advanced Security gives us a lot of options; it can even be set up with Group Policy. And when we take a look at the demo, I'm going to show you some neat things that you can configure.

8.8.2 Configuring an Inbound Rule

Configuring an Inbound Rule
0:00-0:07
In this video, we're going to take a look at managing the Windows Firewall. We're going to open that up. I'm just going to type fire.
Basic Firewall
0:08-0:14
I'm going to take you to the basic firewall, and then we'll look at the advanced firewall. In the basic firewall,
Settings
0:15-0:56
I can turn the Windows Firewall on or off. Notice I have some checkboxes here. You may or may not want to check these. This one would notify you if the firewall wall blocks an application. This one is a little tricky. It blocks all incoming connections, even if it's in the exceptions list. That means there's not going to be any incoming traffic to the server, unless the server initiated it. Very secure, but probably not good to turn on, so don't come in here and say those aren't checked. Let me check them. Bad idea.
I have a firewall for each of my network profiles. Right now, I'm connected to a domain, the domain network profile is in effect. If I actually connect up to a private or a public network, I have a little bit different firewalls for those as well.
Apps to Communicate through Windows Firewall
0:57-1:21
If I need to Allow an application or feature through the firewall, the easiest way to do it is to just click up here. Then we would create an exception. You can see here, we can just create exceptions for applications or Windows features. All I have here is, I can add in an application. I don't have a lot of control over the firewall. If all I'm looking for is to let an application or feature in or out, this is absolutely a wonderful way to do it.
Windows Firewall with Advanced Security
1:22-1:38
If you need to get more detailed than that, what you're looking for is the Windows Firewall with Advanced Security. You can get to it directly from the Start menu, or by clicking Advanced Settings. You can see it opens up a completely different window. Here in the Windows Firewall with Advanced Security, I'm going to get a lot more functionality.
Inbound and Outbound Rules
1:39-1:48
I can create inbound rules that affect traffic that's coming into my computer. I can create outbound rules that's affecting traffic that I might initiate out.
Connection Security Rules
1:49-1:56
Connection Security Rules are used for IPsec. They allow me to specify traffic that's going to be encrypted. Maybe I'm only going to allow it if it's going to be encrypted.
Importing Into Group Policy
1:57-2:06
If I get a firewall set up the way I like it, I can actually export the policy and import it into Group Policy. I can do all of this inside of Group Policy as well.
Logging
2:07-2:32
I can also go into the Properties and set up logging if I need to. I could come in and say, for the Domain Profile, I want to Log any dropped packets. Those are packets that are being blocked by the firewall, or maybe I want to Log successful connections-- things that got through the firewall. If I need to see if my firewall's effective, I can come in and I can do some logging.
We're going to take a look at making an inbound rule.
Creating an Inbound Rule
2:33-2:37
I'm going to right-click and make a New Rule.

Rule Type Tab
2:38-3:15
Right away, I have some choices to make. If I just want to make a rule that affects a program, probably easier to do that in the basic firewall, but I can certainly do that in here. In fact, if I want to control that program just for inbound or outbound, I have to do it in the Windows Firewall with Advanced Security, because when I do it in the basic firewall, it's both inbound and outbound. I could also specify by port to just have a rule that specifies a particular port. We have a bunch of predefined rules, tons and tons of them. If I wanted to go ahead and change one of those based on the rule off it, I can do that.
The most flexibility is going to be in Custom, and there we're going to get all of our choices. I'm going to pick custom and hit Next.

Program Tab
3:16-3:36
First of all, I can say whether it applies to all programs or a particular program. You can see, I could go out and get the EXE and say this is the EXE I'm talking about. I also can specify a rule that applies to services. We didn't see that in the basic firewall. If I want a rule for services, I've got to be in the Advanced Firewall, and I can specify a particular service, all services, whatever I need to do.

Protocol and Ports Tab
3:37-3:56
Once I've chosen my programs and services, I hit Next. I can choose which protocols it applies to. If I select a protocol that supports ports like TCP, then
I can select All Ports, Specific Ports or ranges, Dynamic Ports, and I can select it for both Local ports and Remote ports.

Customize ICMP Settings
3:57-4:24
If I choose some flavor of ICMP and we have two ICMPv4 or ICMPv6, then I get the customized version so that I can come in and say exactly what I'm talking about. Maybe I'm concerned ICMP is going to be used to launch a Deny All service attack against this machine, but I actually want to be able to use Ping to troubleshoot. I come in and say yes, but I'm okay with the Echo Request, that's the Ping, but I don't want any of these other ones to go by. I'm just going to set it to Any Protocol.

Scope Tab
4:25-4:56
Now we can specify which IP addresses this is going to apply to. You can see at the top box, these are local IP addresses. The bottom box are remote IP addresses. I can specify a network range an actual spread of IP addresses or
I can even come in and say look I'm talking about all the Wireless traffic or I'm talking about all the Remote access traffic. I can specify by interface type as well, that way, if somebody adds another wireless card or introduces some more remote access functionality, this rule will apply to that.

Action Tab
4:57-5:21
Once I've chosen my Scope, I set up an Action, and I have three choices. I can Allow the connection, so essentially, I'm making an exception. I can Block the connection, I just don't want that traffic. The middle one says I can have the connection if it's encrypted, so it needs to be encrypted using IPsec. If I choose that, I can come in and customize it. Whether I'm going to require things to be encrypted, I'm going to request that they be encrypted, what sub-protocols I'm going to use.

Users Tab
5:22-5:32
You can also see if I turn it on for IPsec, I get an additional box I didn't have before where I can specify Authorized Users and Exceptions. If I just allow or block, I actually don't get this tab.

Computers Tab
5:33-5:40
I also don't get the next tab, which is the Computers tab that says which authorized computers can connect and any exceptions to this rule.

Profile and Name Tabs
5:41-5:52
Once I've made my choices, I choose which of my three network profiles this is going to apply to, give my rule of name and then Finish.
If I had a policy set up and I exported it.

Control Windows Firewall with Group Policy

5:53-7:10

I made a bunch of rules and I've got to export it; I can also control the Windows Firewall with Group Policy. I'm going to show you that before we go. You go into Tools, Group Policy Management, and we'll just edit the Default Domain Policy. I want to be on the computer configuration because the firewall affects the entire computer. Policies, Windows Settings, Security Settings, and then you can see here, the Windows Firewall with Advanced Security.

If I want to create rules right in Group Policy, I can do it the same way I just showed you. A New Rule, and it's the exact same wizard we just went through.

If I've already exported the policy, I can come in here and import it and just specify that policy I had exported. As a general rule of thumb with Microsoft, anytime you want to make two machines look identical, regardless of what service we're talking about, it's always an export-import. You never want to do the same thing twice, and believe me, if it's me, I can't do the same thing consistently twice.

Summary

7:11-7:21

That's how we work with the Windows Firewall. It's a great, robust firewall. You can certainly go in and do as much as you need to do to make your computer as secure as you need it to be.

8.8.3 Connection Security Rules

Connection Security Rules
0:00-0:15
In this video, we're going to talk about Connection Security Rules. Basically what Connection Security Rules allow you to do is to configure IPsec from within the Windows Firewall with Advanced Security. You also have the ability to monitor any IPsec traffic that's going through.

IPsec
0:16-0:42
What IPsec does is exactly that--encrypt network traffic--and we can choose to encrypt just the headers of the traffic, to make sure that the packet arrived at its destination without being modified. They call that Integrity. It uses a protocol named AH, or we can encrypt the data also known as the pay load, so that nobody can look at the contents of the packet, and that ensures us confidentiality. It uses a sub protocol called ESP.

Connection Security Rules
0:43-0:51
Let's get a good definition of Connection Security Rules, and then I'm going to explain to you how IPsec works, and that should be enough to get you on your way.

Authentication
0:52-1:27
Connection Security Rules have two parts. They involve the authentication of two computers before they begin communications. That's the first piece of it--I need to know that I'm actually having communication with the correct destination computer. There are attacks, called man-in-the-middle attacks, where somebody will set up a computer in between two computers. The end computers believe they're having an encrypted conversation with each other, but in fact, each one has an encrypted conversation with the man-in-the-middle.
With IPsec that's not possible. They authenticate each other before they begin communications. It's mutual communication, and that's done every once in a while throughout the conversation.

Secure Information
1:28-1:41
Once they've done that, they secure the information sent between the two computers by encrypting it with IPsec. If you can remember nothing else from this lecture, if you can remember Connection Security Rules equals IPsec, I would be very happy with you.

Encryption
1:42-1:51
Let's talk for a minute about encryption before we look at how IPsec works. There are two types of encryption: symmetric encryption and asymmetric encryption.

Symmetric Encryption
1:52-2:37
When I think of symmetric encryption, I always think of the movie "The Christmas Story", where a character there gets a Little Orphan Annie decoder ring; he's got a decoder ring and Annie's got a decoder ring. She reads out an encrypted message over the radio, and then Ralphie can decode it. That's exactly how symmetric encryption works. It's the same key on both sides.
If you have a home router, and you go in and put in a passphrase like "I love my router", each of the clients have to put in that same passphrase "I love my router", and then they'll get connected. That's symmetric encryption. The problem with symmetric encryption is how do we deliver that key to the other computer? Not a problem on your home network--piece of cake. You just tell somebody, they put it in.

Asymmetric Encryption
2:38-3:14
When we're talking about a business network, it may not be convenient to send that phrase out over the network. We use something called asymmetric encryption, where we have one key that encrypts, and a different key that decrypts. I can give anybody my public key-- the one that encrypts--because all that they can do with it is encrypt something and send it to me. I keep my private key very private, meaning that I'm the only one that can decrypt that message. A lot of people are very surprised when they hear that symmetric encryption is the better type of encryption. The reason is, we tend to use that for session keys, and so they're swapped out very quickly.

How IPsec Works
3:15-3:44
Here, we're going to take a look at a client communicating with a server. It could two clients, it could be two servers, it makes no difference. What we want to do is exchange those session keys, which you're going use symmetric encryption, but before I can do that, I have to come up with a secure way to exchange them, so IPsec goes through a process that we call Main Mode Security Association. Any time you're talking about IPsec, Security Association pretty much means the key. It's the key, and the process of getting the key.

Main Mode SA
3:45-3:54
With main mode, this is the time when the two computers authenticate each other, and they establish an initial encrypted channel that they can use to exchange session keys called the Main Mode SA.

Kerberos
3:55-4:14
How can they do this? There's three ways, one way is they might use Kerberos. Kerberos is the Active Directory Security Protocol, so if these computers are in the same domain or forest with each other, they can use Active Directory and it's no problem. If they want to interact with UNIX or Linux and that supports Kerberos, that would work just fine too.

Certificate
4:15-4:26
If I can't use Kerberos, then I can use a certificate, which is that public and private key that I just talked about. The only problem is I'm going to have to get these computers certificates, which may or may not be quite a bit of work.

Pre-Shared Key
4:27-4:50
The third thing I can use is exactly what you use on your home network, and we call that a pre-shared key, because if your password for your home router is "I love wireless", you're going to have to pre-share that with me so I can type it in on my computer and I can get connected.
Using one of these three methods, it does the Main Mode SA, it establishes an encrypted channel, and now it's going to exchange two session keys.

Quick Mode SA
4:51-5:07
These session keys are called Quick Mode SAs, so there'll be one going that direction, and another one going in that direction. We know these keys are secure because we had our Main Mode SA when we set them up, and now we're very secure, because we're using different keys depending on the direction, and we're extremely happy.

Summary
5:08-5:55
Connection Security Rules allow me to specify different circumstances under which traffic must be encrypted, and then if I'm requiring IPsec, I can also set up how Main Mode encryption is going to go. I can use Kerberos, I can use certificates, I can use a pre-shared key, I can specify which pieces of the packet will be encrypted.
It supports something called Domain Isolation, where I can isolate computers in my domain simply by requiring that they use IPsec to communicate; that way, anybody who's not in the domain won't be able to communicate with those computers. In order for two computers to use IPsec, they have to both have a policy. So when you create your Connection Security Rules, make sure that they apply to both computers that are going to be part of the conversation. The best way to do this is to set up your Connection Security Rules in Group Policy.

8.8.4 Authenticated Firewall Exceptions

Authenticated Firewall Exceptions
0:00-0:08
In this video, we're going to talk about authenticated firewall exception. If the computer can authenticate, it will be exempt from the firewall polices.

Exemption Because of IPSec
0:09-0:32
Any computer or user who connects using IPsec is authenticated, because part of IPsec is for the computers to authenticate each other. If I can authenticate and have a connection with IPsec, then we have a pretty good relationship. So I can go ahead and make that computer exempt from a particular firewall policy or rule, and allow it to access all the ports and services on the server just by virtue of the fact that it's authenticated and it's using IPsec.

Using Group Policy
0:33-1:02
If I want to do this in Group Policy, I'll go into the Computer Configuration, Administrative Templates, Network, Network Connections, and then there's an options for the Windows Firewall. When I click on that, I'm going to see a setting that says "Allow authenticated IPsec bypass". When I turn that on, I'll see a textbox that says "Define IPsec peers to be exempted from the firewall policy". What I have to do is type in an SDDL, which is like a security descriptor; it's a string that corresponds to the groups for the computers to which this policy applies.

Enable in the Firewall
1:03-1:50
If I don't want to do this with Group Policy, I want to do it in the individual firewall. I can create a new Inbound Rule, specify the traffic to which the rule applies. It could be all traffic, it could be just certain traffic. I'm talking about ports, protocol, that type of thing. Then when I get to Action, I'm going to say Allow the connection if it is secure, which is code for "it's only going to be allowed if it uses IPsec."
There's a setting inside of here that you can customize that says "Override block rules". Essentially, if the client connects using IPsec, it's exempt from any firewall block rules. Now once I specify "Allow the connection if it's secure", I'm going to get two new pages in the wizard. One for Users, the other for Computers. I would select "Only allow connections from these users or computers" on the appropriate page-- specify the groups that I'm going to allow to override the policy.

netsh Command
1:51-2:38
Here's an example if you want to do this at command prompt using netsh. We have netsh. We're working with the Advanced Firewall, and going into the firewall, I'm adding a rule. The name of my rule is Inbound Secure Bypass Rule. It's an inbound rule. I'm going to require authentication, and if they authenticate, I'm going to let them bypass. This is the part where it gets tricky. This is specifying the groups for computers that are allowed to bypass. This whole string here is that SDDL. We've got a few acronyms up here, but basically you've got to get the SID--the actual SID of the group--so, from this S all the way down to the 4. That's the SID of a particular group in Windows. I can also specify remote user group again, the same thing, but I've got to supply the SID of that group. But this would all be one command at the command prompt.

Summary
2:39-2:53
Authenticated Firewall Exceptions really just mean that anybody who authenticates and communicates with the server using IPsec is presumed to be safe, and therefore, we can override some of the block rules that are in place. That's really all Authenticated Exceptions does.

8.8.5 Firewall Policy Facts

Windows Firewall with Advanced Security is software that provides real-time protection from unwanted access such as hackers, viruses, and worms. By default, all outbound traffic is allowed (as are inbound responses to those requests), and all unsolicited, incoming traffic is blocked.

Use Group Policy to manage Windows Firewall with Advanced Security.

Feature	Description
Profiles	Windows Firewall with Advanced Security uses *profiles* to group settings (i.e., firewall rules or connection security rules) according to the network where the computer may be connected. The Network Location Awareness Service determines the type of network connection and Windows Firewall applies the settings for the connection you've made. There are three types of profiles: • The *domain profile* specifies the behavior for when a computer is connected to its corporate domain. • The *private profile* specifies the behavior for when a computer is connected a to private network location. • The *public profile* specifies the behavior for when a computer is connected to a public network location. In Windows Server 2008, only one profile is applied at a time (the most restrictive profile), even with multiple network adapters. Beginning with Windows Server 2008 R2, each network adapter applies the firewall profile best suited for the connected network.
Firewall rules	Firewall rules are applied to either inbound or outbound traffic: • *Inbound rules* block or allow inbound traffic that matches the rule criteria. By default, inbound traffic is blocked when Windows is installed. You must create inbound rules to allow inbound traffic. • *Outbound rules* block or allow outbound traffic that originates from a computer that matches the criteria in the rule. By default, outbound traffic is not blocked. You must create outbound rules to block outbound traffic. The criteria for the rules are: • **Program** rules control connections for a program. • **Port** rules control connections for a TCP or UDP port. • **Predefined** rules control connections for a predefined Windows experience. • **Custom** rules are defined by the administrator or user creating the firewall rules.

	Each incoming packet is inspected and compared to criteria in the firewall rule. If the packet matches the rule, the specified action is taken: allow the connection, block the connection, or allow the connection if it meets specified criteria. Creating a firewall rule to allow traffic does not secure that traffic. You must use connection security rules to secure the traffic.
Connection Security rules	Connection security rules ensure that connections between two computers are authenticated or encrypted. Windows Firewall with Advanced Security uses IPsec to secure traffic in transit over the network. Connection security rules require that both communicating computers have matching connection security rules or an IPsec policy. You have to create a firewall rule to allow network traffic protected by a connection security rule.
Policy file	A *policy* is the overall combination of your Windows Firewall with Advanced Security settings that you have exported to a policy file. • You can export a policy configured on one system and import it on another system. • An imported policy overwrites and is applied in place of the current policy. • Policies are saved with the **.wfw** extension. • You can import policies into GPOs to apply those policies to multiple computers.

8.9 Application Restriction Policies

As you study this section, answer the following questions:

- How are software restriction policies managed?
- What type of restrictions can be implemented using software restriction policies?
- How does a hash rule identify one application from another? Does the application have the same hash value after an update is applied?
- What are the advantages of AppLocker policies over software restriction policies?
- Which file extensions can an AppLocker script rule be applied to?
- If software restriction policies and AppLocker are configured on the same object, which takes precedence?

After finishing this section, you should be able to complete the following tasks:

- Configure a software restriction policy for a specific user.
- Create a path rule for an application.
- Create a hash rule to create software restrictions.
- Create an AppLocker executable rule.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 6.0 Group Policy.
  Configure Application Restriction Policies

This section covers the following 70-410 exam objective:

- 603 Configure application restriction policies.
  This objective may include but is not limited to:
  Configure rule enforcement
  Configure AppLocker rules
  Configure Software Restriction Policies

8.9.1 Software Restriction Policies and AppLocker

Software Restriction Policies and AppLocker
0:00-0:23
All right, we're going to talk about Software Restriction Policies and Application Control Policies, which are also known as AppLocker. Group Policy is great for installing software, but for a long time it didn't have any ability to block software. Software Restriction Policies allow you to do just that. I can go in and I can specify applications that I either want to run in my environment or don't want to run in my environment.

Software Restriction Policies
0:24-0:29
Basically under Software Restriction Polices, we have four different types of rules that we can create.

Hash Rule
0:30-1:17
The first type of rule would be something called a hash rule. So let's say I have a problem in my network with users bringing in peer-to-peer file sharing software, and I don't want that being installed. We'll call it "FileShare1". I go out and I get the EXE for FileShare1; I specify it in the hash rule. The computer makes a hash out of the EXE, which is like an electronic fingerprint.
Now, it doesn't matter what they would call that EXE. If the computer sees that EXE, it's going to ban it. The problem with the hash rule is it's way too narrow. So if they find out that I've banned FileShare1.0, they'll go and get FileShare1.0.1, I'll ban 1.0.1, they'll get 1.0.2. I'm going to find myself chasing different versions of this program in order to ban it.

Certificate Rule
1:18-1:44
A second type of rule I can create is a certificate rule, so maybe I decide to go out and get the certificate that the company uses to digitally sign their software, and I can ban anything that's signed by that particular manufacturer. Let's say, for the sake of example, that FileShare1.0 is put out by Microsoft; well, I get the Microsoft certificate, and now I ban everything Microsoft. That type of rule is sometimes too wide a net.

Path Rule
1:45-2:15
Next type of rule you can create is called a path rule, and in a path rule, I would specify a particular path on the hard drive or in the registry, and basically say anything that installs itself into Program Files--FileShare1 or FileShare--I don't want that.
Most applications give the user the opportunity to install that software into another place. If they figure out that that location is banned, they'll just install into C:\boringworkstuff, and then they'll be good to go.

Network Zone Rule
2:16-3:13
The last type of rule that I could create with Software Restriction Polices is a Network Zone Rule. That would basically say, anything that is coming from a particular zone Internet Zone, Intranet Zone, Local Computer Zone, Restricted Sites, Accepted Sites. There's five different zones that I can specify, and anything coming in from those zones would either be allowed or denied. Software Restriction Polices--this is the original way they had for us to restrict software--and this is going to be for everything pre Windows 2008 R2 or Windows 7, because you're going to see once we get into R2 and Windows 7, you'll have a lot better options. Here's my Hash Rule, which is the electronic fingerprint--a little too narrow. Certificate Rule, which gets the certificate of the manufacturer and bans that--a little too broad. Path Rule that bans a specific path, very easy to get around, and Network Zone Rule, where I can ban it by network zone.
It's a great idea to block software, but a little bit too rudimentary until we get into 2008 R2 and Windows 7.

AppLocker
3:14-3:47
In Windows Server 2008 R2 and Windows 7, Microsoft came up with something new, and they're called Application Control Polices, also known as AppLocker. Much better functionality. In Software Restriction Polices, I can simply say, "This software is allowed or denied," I could white list it or black list it, but it would affect everybody, even administrators.
With Application Control Polices, I can specify particular groups that these rules will apply to. I can make rules for Executable Files, make rules for Windows Installer Files, MSI files, and rules for scripts.

Custom Publisher Rule

3:48-4:44

The real power with AppLocker comes in with the Custom Publisher Rule. If I do a Publisher Rule, I can go through and specify an application and have it ban all versions of that application, so I can go out and get the EXE for FileShare1.0 and say hey, I want to ban FileShare1.0 and everything above.

I can also configure exceptions to that rule. Let's say my company is standardized in Office 2010, and Office 2013 is out, and I've noticed that some people start bringing that in from home and installing it. Well, that's a licensing problem, that's their home version of Office, and it shouldn't be installed in the company. I can go in and create a rule that bans all versions of Microsoft Office, and then on the Exceptions tab, I can include an exception that will only allow 2010. I can do that for each of the Office applications: Word, Excel, PowerPoint, and so on, but that's the great thing--I can have it ban all the versions and then have an exception.

When to Use Software Restriction Policies

4:45-5:24

AppLocker is much more flexible than Software Restrictions Policies. The only reason we would go with Software Restriction Policies is if we have older clients, so if anything is below Windows 7, i.e., Vista or XP, then I have to do a Software Restriction Policy. As long as I'm at Windows 7 and Windows 8, I should be good to go. Application Control Polices-- they're also known as AppLocker, but I didn't write that down--came in with Windows Server 2008 R2 and Windows 7.

The powerful one is the Custom Publisher Rule, which allows me to specify versions. I can even have exceptions, and I can specify the groups to which this rule will apply. Now as you'll see in the demo, these things are really cool to set up.

Default Rules

5:25-5:45

There's only a couple of things that you should keep in mind. The first thing is, you want to make sure that you create the default rules. The default rules basically say that the operating system is allowed to run. If you don't create the default rules on a computer that's in a work group, you have to reinstall the operating system. If it's in the domain, you'll have to change your policy and have them reboot their computer.

Enforce the Rule

5:46-6:12

The second thing that you're going to need to make sure that you do is that you enforce the rule. By default it's not set up to enforce the rules, and the reason for that is this: maybe I don't want to enforce the rules I just want to audit, so there is a selection in there that we'll see in the demo where I can choose to audit the rules, maybe run it for a week, see what would be blocked, and then when I'm happy, I can turn on Enforcement. If I don't select Enforce, then the applications are not going to be blocked.

Application Identity Service

6:13-6:26

The last thing I need to remember about AppLocker is that it requires the Application Identity Service, so you're going to need to make sure that the Application Identity Service is set to Automatic and running on all the clients, and that can be done with Group Policy.

Summary

6:27-6:54

Software Restriction Policies and AppLocker, or Application Control Polices, both intended to block software. Software Restriction Policies, kind of rudimentary, but all that we had available up until Windows Server 2008 R2 and Windows 7. Application Control Polices, also known as AppLocker--much better, a lot more fun. We can actually specify version numbers, exceptions, groups--just need to make sure that we enforce it and start the Application Identify Service.

8.9.2 Configuring Software Restriction Policies

Configuring Software Restriction Policies
0:00-1:04
In this video, we're going to take a look at configuring Software Restriction Policies. Software Restriction Policies allow you to allow or prevent software running on the computer. This is what you have to use if you've got clients that are not Windows 7 or Windows 8, because with Windows 7, we get AppLocker, which is a lot more efficient. Let's go in and take a look at it. We need to get into Group Policy. I'm going to go in through the Tools menu. Software Restriction Policies are available in all of the Group Policies: Local Policy, Domain, any of the ones you make up at the domain level, OU, Sites. We'll just edit the Default Domain Policy because we're interested in taking a look at how they work. We're going to open up Policies, Windows Settings, Security Settings, and then here are Software Restriction Policies. You can see by default, there are no policies that are defined.

Creating a Software Restriction Policies
1:05-1:09
To start out, we're just going to right click and make a New Software Restriction Policies.

Security Levels
1:10-1:55
Security levels are where you start from. We tend to refer to them as "blacklisting" or "whitelisting." Right now, you can see that the checkmark is on Unrestricted, which means that all software is allowed to run on the computer, except those for which I specifically make a rule denying it. That's going to be easier to support than Disallowed, where everything is disallowed unless I make a rule to allow it to run. You'll see in a minute why that's pretty tough. You'd have to go through and find every single thing that needs to be allowed to run and that would require a lot of testing. Unrestricted is a little bit easier to use because you can ban the software that you don't want.
When we go into Additional Rules, you're going to see that two rules exist by default.

Additional Rules
1:56-2:31
The first rule here is talking about the System Root, the next rule is talking about Program Files Directory. Basically, what it says is that anything that's in System Root or that's already installed in Program Files is allowed. These are the default rules. It's not good to get rid of these because essentially, you can prevent Windows from functioning. Don't mess around with the default rules. Leave those there. Just make additional rules to support whatever you need to do.
We have Enforcement, which adds a little bit of extra restrictions or not, depending upon what you want.

Enforcement
2:32-3:36
You can see that we can apply the software restriction rules to all of them except Libraries, like DLLs, or we could do All files. It would be very difficult to ban all the DLLs; you'd have to find out exactly which ones need to run. You can see by default, it's everything except these system files. If you did need to affect DLLs, you certainly could adjust the Enforcement Properties.
We can set up to whom we're going to apply software restriction policies, all users or all users except local administrators. That's really as far as we have in terms of functionality specifying by user. You'll see when you take a look at Application Control Policies, it changed that a little bit. We can also choose to Enforce certificate rules or Ignore certificate rules. You can see that Microsoft even says "Certificate rules will negatively impact the performance of your machine." By default, they're set up to ignore them. We can go in here and adjust what's considered to be executable code, and add any extensions that we need to add in, or remove any that we don't want to be considered executable as well.

Designated File Types
3:37-3:52
These are fairly basic.

Create Additional Rules
3:53-4:11
What you do is you come in and you create additional rules to specify the software that you're trying to restrict. When we create our additional rules, there are four types of rules that we can use. Basically, they're enforced in the order that I'm going to go through them.

Hash Rule

4:12-5:29

The first one we'll take a look at is a Hash Rule. A hash in this context-- they tend to refer to it as a thumb print for the software. Basically, what you do is you go out and you find the EXE or the software that you're looking to control.

I'm just going to select Notepad because it's something I know where it is. All we're looking for is to take a look at it. I find the EXE that I'm looking to control. Then all I can do is say whether it's Disallowed, Unrestricted, or it goes with the Basic User policies. I'm just matching it up to my Enforcement Settings. What the hash rule does is it makes a hash, which is a unique number based off of that EXE. That's why we'll refer to it like a fingerprint or something like that. The computer will know this EXE if it runs. Even if it was named something different, it can uniquely identify this particular software.

I tend to think of a hash rule as a little bit too narrow a net, because it's only going to catch that specific application, and sometimes that's not what I'm looking for. We say, "Maybe there's something better I can try?

Certificate Rule

5:30-6:18

Maybe I could try a Certificate Rule?" We've already seen they negatively impact the performance of the machine. There's a little bit of other problems with these things as well.

What I would need to do is go out and get the certificate that the manufacturer uses to sign their software. Then I would put that certificate in here, and essentially at that point, what I'm saying is every single bit of software by that manufacturer is allowed or disallowed. Where hash rules are a little bit too narrow a net, this can sometimes be too broad, depending on how much software that manufacturer makes and whether it's present in your environment. We're saying, maybe a Certificate Rule is not exactly what we're looking for. Let's see if we can do something else.

Path Rule

6:19-6:54

Third one is a Path Rule. In here, this would be a path. It actually could be in the file system or the registry, but we would be saying anything that's in that path is disallowed. That will prevent them from running it or installing it if it goes into that particular path. The registry ones--if you can find a registry key that you can ban that's going to keep it from running--great. File paths--not very effective, because most software will have a spot in the install saying, "where do you want me to put this in the file system?" This is a net that's too wide.

Network Zone Rule

6:55-7:37

Let's look at our fourth option. That would be the Network Zone Rule. This is going to ban anything coming in or allow anything coming in from that particular network zone. These are based on the ones that are in Explorer. I can allow or ban anything coming from the Internet, Local intranet, Restricted sites, Trusted sites; those are the four zones that we see in Internet Explorer. I can also ban or allow anything coming from the local computer. Again, little bit too wide of a net, because there may be things that I do want them to be able to install from the local computer. It's just a particular software I'm looking to get rid of.

Summary

7:38-8:16

Software Restriction Policies really were great when they came in. The only problem with them is that it was difficult to specify exactly what it was you wanted to restrict without it becoming a lot of overhead to continue to support that. If you're looking for something more flexible--you've got Windows 7, you've got Windows 8--what you're really looking for are Application Control Policies. If you have older clients and you need to ban a particular piece of software, then for what they do, at least they give you some control over that, and you can go ahead and enforce company policy on the workstations. That's how we use Software Restriction Policies.

8.9.3 Software Restriction Policy Facts

Software Restriction Policies allow an organization to control the applications that run on the computers in the environment. You can use software restrictions to:

- Identify software that will be allowed to run on a computer.
- Restrict the programs that users can run on a shared computer.
- Determine software packages that can be installed on a computer.
- Run only digitally signed scripts.

AppLocker, introduced with Windows Server 2008 R2 and Windows 7, is a more robust tool for controlling software execution. Microsoft recommends that you use AppLocker for Windows 7 and later.

By default, there are no Software Restriction Policies configured. Executable files run based on the permissions that user or groups have in the NTFS file system.

Software restriction can be implemented in the following ways:

Type	Restrictions
Unrestricted	All applications are allows to run, except those specifically excluded.
Disallowed	All applications are prohibited from running, except those specifically allowed.
Basic User	All applications that standard users can run are allowed. Applications that require administrative privilege are not allowed to run.

Once the type of restriction is set, you configure the software restriction rules. The software restriction rules specify the conditions under which applications are allowed or denied to run. The following table identifies these rules:

Condition	Description
Hash	The *hash* condition uses the digital fingerprint (also known as a *file hash*) of the application. • A hash value of a file is based on the content of the file, not the name of the file. • You must recreate file hashes each time the software is updated or changes version.

Certificate	The *certificate* condition uses the digital signature of the application's publisher. The digital signature contains details about the company that created the application. • The digital signature is extracted from the application file. • If the file does not have a digital signature, it cannot be used with the publisher condition. • The certificate condition applies to all applications from the specified publisher.
Network zone	The network zone condition specifies where the application originated. Options include: • Internet zone • Intranet zone • Restricted sites • Trusted sites • Local computer zone
Path	The *path* condition specifies a folder, a file, or a wildcard of files to restrict or allow execution. • If you specify a folder, restrictions apply to all programs within that folder. • Path conditions are the least secure of all the software restriction conditions. • When using this condition, implement NTFS permissions to prevent users from copying executable files to locations outside the scope of the path condition.

8.9.4 Configuring AppLocker

Configuring AppLocker
0:00-0:58
In this video, we're going to take a look at Application Control Policies, also known as AppLocker. AppLocker can be implemented in any Group Policy, Local Group Policies, or the ones that are stored in Active Directory.
We're going to go into Group Policy Management Console. I'm going to get in through the Tools menu, but you could also use the Start menu and, because we're just interested in looking at this, I'm just going to edit the Default Domain Policy so we can take a look. You want to go ahead and expand Policies, Windows Settings, Security Settings, and then down here we you can see Application Control Policies, then you have to go into AppLocker. You'll see those terms used interchangeably in documentation and exams. Some people say Application Control Policies, some say AppLocker. Either one is accurate.
We have different types of rules we can create.

AppLocker Rules
0:59-1:22
Executable rules apply to .exe .com files. Windows Installer rules would be Windows Installer files like .msi. Script rules would be anything that's a script--.bat, .js1, a whole bunch of extensions there, and then Packaged App rules apply to Packaged Applications, which would be a .appx. You've got to start out knowing what type of file you're talking about.

Creating a Rule
1:23-2:06
I can create my file a number of ways. I can go in and Create a New Rule. I could Automatically Generate Rules, or I could give it a folder, and it would find all the applications inside that folder and give me suggested rules so that I can then get rid of them or add to them or do whatever I want, or I can create the Default Rules. If I don't create the default rules, at the end of creating a new rule, it will prompt me to create the default rules. I'm not going to be proactive. I'm just going to jump in as if I'm creating a new rule. It says, before you get going, make sure you have the application installed. Make sure you've got everything backed up.

AppLocker vs Software Restriction Policies
2:07-2:39
AppLocker was intended to be an improvement on the old Software Restriction Policies. Right off the bat, we can see that we're getting additional functionality over those. The Software Restriction Policies just allow me to allow or deny an application, but that's for all users. Here, I can actually have my AppLocker rule affect a particular group or even a specific user. I've already got some additional functionality in here if I need this rule just to apply to a particular group or an actual user.

Path Rule
2:40-2:49
The Path Rule is just like a Software Restriction Policy rule in that it would block everything in a specific folder, or you can specify a path in the registry.

File Hash
2:50-3:03
File Hash would be for an application that's not signed. It would provide just that particular .exe. It makes a hash, which is a unique identifier. They'll talk about it as a fingerprint of the file, and it will ban that particular file.

Publisher Rule
3:04-3:38
The types of rules that AppLocker is famous for are the Publisher rules, because these give you really the most flexibility in identifying the software that you're trying to allow or block.
We start creating our Publisher rule by going out and getting the .exe for the file that we want to control, and I'm just going to go into Windows System 32 and get Notepad, honestly, because I know where it lives. We open that up.
As the rule stands right now, I would just be blocking Notepad.

Use Custom Values

3:39-4:52

Here's the power of the Publisher rule. With the Publisher rule, I can actually go in and say, "Use custom values", and then that lets me go in and adjust any of these settings.

Now, with the hash rule, basically what happens is this: You would go out and get the .exe and ban that particular version. Let's say you're having a problem with Notepad. It wouldn't really be Notepad, but I'm just using an example, and you've gone out and you've got this file version 6.2 blah, blah, blah. In the hash rule, we're just banning that particular .exe. What the user's going to do is go out and get Notepad 6.3, or 6.4, or 5.2.

The great thing about this is I can go in and say, "All right. I'm talking about Notepad 1.0.0 and above". Now, I've banned all versions of that application. If I had a problem, let's say, with a peer-to-peer file sharing software named "Fileshare", and I want to make sure that when the new versions of Fileshare come out, this rule's going to affect it. I can go through and make a Publisher rule and ban all versions and above. I could also ban anything with that file name, product name, publisher. I can adjust it a little bit.

Create Exceptions

4:53-6:39

The other great thing about AppLocker is I can create exceptions. What commonly happens is, your company may have purchased a particular version of a software, and that's okay in the environment, but what you don't want is users bringing in a different version of that software. I can give you a great example.

We were standardized on Office 97. Outlook 98 came out. We had a lot of executives begin bringing that in, installing it on their workstations, which is illegal, because we didn't own any licenses for it, but even worse than being illegal--if there's anything worse--it would break the link between anything Outlook 98 and anything Outlook 97. If that person had an administrative assistant, the administrative assistant could no longer manage their calendars, send email on behalf of, and even better, you couldn't uninstall Outlook 98. You had to re-image the machine.

If AppLocker had been available at that time, what I could have done is banned all versions of Outlook from whatever the first version is and up. Then, on the next page, I can make an exception, so I can say, Yes. It's all versions of Outlook or all versions of Notepad 1.0 and above, except for the one I want them to be able to use. Now, I've gone through and said, Yes. They can have this particular version, but that's it.

Now they can't upgrade it. They can't use a different one. They're locked into the one for which I actually have purchased the software. The company's not exposed to any issues with piracy, and I don't have any problem supporting versions that haven't been tested and released into production. Then, I could go through and I could create my rule.

Default Rules

6:40-7:59

If you don't create the default rules before you make your first rule, it comes up with this great dialogue box saying, essentially, you really need to create the default rules. Make sure that you click Yes. Very, very important. The default rules say, everything that's in Program Files, everything that's in Windows, anything that's being run by an administrator, all of these things are allowed to run. Here's my Notepad rule, but since I didn't change it from Allow to Deny, it says Allow.

If you do not create the default rules, and that policy hits a Windows 7 or Windows 8 machine, that machine is done. On a standalone machine, if you're working in the Local Group Policy, that's actually a re-image.

In a domain computer or something that's showing to Active Directory, you'd have to come in, redesign your AppLocker rules, and, because it's done in the computer half of the policy, you'd need to get them to reboot their machines. At least in Active Directory environment, you can fix that mistake. Don't make the mistake to begin with. Make sure that you either proactively create the default rules or you agree to create them at the end of the wizard. AppLocker is fairly simple to support, but there's a few things that you've got to keep in mind because, as it stands, I've created rules, nothing's going to happen.

Rule Enforcement

8:00-8:19

There's two things that I have to do before AppLocker will take effect. The first is to set up Rule Enforcement. You can see, by default, none of the rules are enforced. I can make all the rules I want. Nothing's going to happen until I come here and I say, Yes. I want to enforce these types of rules.

Enforce Rules or Audit Only

8:20-9:05

We have two choices: Enforce rules or Audit only. Audit only will not block the application. It will just send a message to the Event Log every time that application would have been blocked, which lets you run this policy and test it for a little while, make sure it really is working the way you want it to work, instead of rolling it out and then getting a bunch of tickets that say, hey, I can't run the correct version of Notepad because the rule was created incorrectly. You can start with an audit. Make sure everything's going to be correct before you actually set it up to enforce it. Make sure that you enforce the rules. If you don't enforce the rules, nothing's going to happen. That's the first extra thing that we need to do that the software doesn't really prompt you to do.

Application Identity Service: Set to Automatic and Start

9:06-10:11

The second thing that needs to be done is that a particular service needs to be started and set to Automatic, and I'm going to show you this service inside of the Services Console, but you can turn this on via Group Policy.

In order for AppLocker or Application Control Policies to function, the Application Identity Service must be started and set to Automatic. I would come in here, set it to Automatic, and then Start it. Because it's set to Automatic, it's always going to start when the computer reboots and now, it'll be able to identify the applications and block it. Make sure that you know it's the Application Identity Service. A lot of people just hear the word application. There's four or five services here that all start with Application. It's the Application Identify Service that AppLocker requires.

Using Group Policy

10:12-10:39

In real life I would not visit every workstation and turn this on. What I would rather do is turn it on via Group Policy. Most likely in the exact same AppLocker Policy, I would go into Systems Services, Application Identity, and set that up to start automatically; and then, when the computers reboot, they're going to get both the AppLocker Policy and the Application Identify Service will be started.

Summary

10:40-11:21

AppLocker, awesome functionality. Really fantastic. I can go through. I can specify all versions of a program. I can create exceptions that only the particular version that I'm interested in is going to be able to run. Make sure that you know the different types of rules that you can make. Executable rules apply to executables, whereas Windows Installer Rules apply to MSI, and then make sure you know that the rule that gives me all this great functionality is that Publisher Rule. That's the one that's really important.

That's how we set up AppLocker so that we can control what types of software can run on the workstation.

8.9.5 AppLocker Facts

AppLocker policies (also known as *application control policies*) were introduced with Windows 7 and Windows Server 2008 R2. AppLocker policies are similar to software restriction policies, but have the following advantages:

- The Automatically Generate Rules Wizard reads the contents of a specified folder and generates recommended rules based on the folder contents.
- Polices can be applied to a specific user or group.
- AppLocker provides flexibility in identifying the software to allow or block. Options include all products from a publisher, all products with a specified product name, and all files with the specified file name.
- Policies can be applied to all existing, future, or previous versions of an application.
- Exceptions can be included in policies.

The following table describes the AppLocker rule types. For each rule type, you specify the users to whom the rule applies and the conditions for applying the rule.

Rule Type	Description
Executable	An *executable* rule applies to files with .exe and .com extensions. When you create a rule, the scope of the rule is set to Everyone. If you choose to modify the rule, you can select a specific security group or user account.
Windows Installer	The *Windows installer* rule applies to .msi and .msp file extensions. You can control the installation of: • Installer files based on whether the files have a digital signature. • Installer files based on user. You can also combine user with the digital signature requirement. For example, only Administrator can install an .msi file without a digital signature. • Software or software updates through Group Policy.
Script	The *script* rule applies to .ps1, .bat, .cmd, .vbs, and .js file extensions.
Packaged app	*Packaged app* (.appx) rules apply to Windows applications that are purchased through the Windows Store and can be used on devices running Windows 8, Window 8 RT, and Windows Server 2012. • All of the executable files, Windows installer files, and scripts for Windows packaged apps have the same identity (software publisher name, product name, product version, etc.). • Packaged apps can be controlled in AppLocker with just one rule using the single identity.

When you create a new rule, you must specify a condition for the rule regardless of the rule type. AppLocker uses *conditions* based on file properties to enforce rules. AppLocker rules have the following conditions:

Condition	Description
Publisher	The *publisher* condition uses the digital signature of the application's publisher. The digital signature contains details about the company that created the application. • The publisher condition provides the greatest flexibility in applying AppLocker restrictions. • The digital signature is extracted from the application file. • If the file does not have a digital signature, it cannot be used with the publisher condition. • You can use the slider on the left of the publisher options to specify Any publisher All products, files and product versions from a publisher. A specific product, files name and file version for a specified publisher. All versions, a specific version, previous versions, or future versions of a product. Any combination of publisher, product name, file name, and file version. • The **Use custom values** option gives you additional flexibility in specifying the publisher settings.
Path	The *path* condition specifies a folder, a file, or a wildcard of files to restrict or allow execution. • If you specify a folder, restrictions apply to all programs within that folder. • Path conditions are the least secure of all the AppLocker conditions. • Implement NTFS permissions to prevent users from copying executable files to locations outside the scope of the path condition.
Hash	The *hash* condition uses the digital fingerprint (also known as a *file hash*) of the application. • A hash value of a file is based on the content of the file, not the name of the file. • You must recreate file hashes each time the software is updated or changes versions.

Be aware of the following:

- The Application Identity service (AppIDSvc) must be started and running on the client for AppLocker rules to be enforced. You can set the service to start automatically using a Group Policy. Beginning with Windows 7 clients, AppLocker rules take precedence over software restriction policies.
- If both software restriction policies and AppLocker policies are configured on the same policy object, only the AppLocker settings will apply on computers using Windows 7 or later. Microsoft recommends that you use AppLocker and not software restriction policies for Windows 7 and later.
- If no rules have been defined for a specific type, then all applications of that type are allowed to run. Once you define a rule, then only software allowed by that rule (or the default rules) is allowed.
- Exceptions allow you to specify a condition that is exempt from the AppLocker rules.
- In order for AppLocker to take effect, you must configure rule enforcement.
- AppLocker has a *soft-enforcement* (also known as *auditing*) mode. Soft-enforcement mode:
  - Uses restrictions to only monitor AppLocker events. Blocked software is still allowed to run while in soft-enforcement mode.
  - Audits AppLocker functionality before full implementation in the environment.
  - Verifies which applications are affected without actually blocking or hard-enforcing the applications from executing.

The enforcement mode (either **Enforce rules** or **Audit only**) applies to all rules of a specific type. You cannot selectively enforce or audit different rules within a rule type. For example, you cannot audit one executable rule and enforce another executable rule, but you can audit all executable rules and enforce all script rules.

- Events that are generated by auditing AppLocker are written to the AppLocker event log. Each log contains the following information:
  - Rule name
  - SID of the user or group
  - File and path of the restricted or permitted application
  - Rule type or condition used
- In Windows 8 and Windows Server 2012, each file within a classic desktop application can have a unique identity including software publisher name, product name, product version, etc. Each of these individual components must be controlled separately within AppLocker. app

You must enable the default rules for each rule type. For example, failure to specify a default executable rule prevents Windows from executing system files in the **C:\Windows** and **C:\Program Files** directories.

8.9.10 Group Policy Preferences

As you study this section, answer the following questions:

- How do Group Policy preferences allow you to centrally manage settings on the workstation?
- How would you apply Group Policy preferences to pre-Windows 7 clients?
- Which Group Policy preferences would be the most useful for typical networked Windows systems?

After finishing this section, you should be able to complete the following task:

- Configure group policy preferences to centrally manage settings in the user environment.

8.10.1 Group Policy Preferences

Group Policy Preferences
0:00-0:17
Let's take a step back and review what we know about Group Policy. If we're going to do something that's a regular function of the computer that's going to be in Control Panel. Group Policy is really for advanced settings that Microsoft wants administrators to control. Regular users don't usually know how to get into Group Policy.

Group Policy, Registry Edit, and Control Panel
0:18-0:23
Then, anything that's really obscure, they don't want you to change very often at all--that would be done with a Registry Edit.

Logon Scripts
0:24-1:07
Now that we have Group Policy at the domain level, which came in with Active Directory, that gives us our Centralized Administration which we're always looking for. But before Active Directory existed, before Group Policy, any type of central management would be done with a logon script, and even after Group Policy came into being with Windows Server 2000, there remain certain functions that could only be done with logon scripts. Even till today, you might find companies where they've got Window Server 2008 or 2008 R2; they've got Group Policy set up, but they're still doing some things with logon scripts, like mapping drives, adding environmental variables to the computer, creating local users, maybe even dropping a file on the desktop, or adding something to the registry.

Group Policy Preferences
1:08-1:51
In Windows Server 2008 R2, Microsoft added a whole category to Group Policy-- both on the computer side and the user side--called Preferences, and basically the idea is to take these things which still, up until that version, had to be done in logon scripts, and move them into Group Policy. With Group Policy Preferences, I can do things like that. I can map drives based on the OU or the group that the user's in. I could even create a user on the local computer, maybe as a backup user, in case the computer can't talk to the domain. I can set environmental variables. I could put a shortcut on their desktop. Maybe I want to put a shortcut to the intranet website on every desktop in the company; I can totally do that.

Client Side Extensions
1:52-2:47
One thing I want you to keep in mind is this though: Group Policy Preferences came in with Window Server 2008 R2. That means they will only be processed by Windows 7 clients or Windows 8 clients. Since they didn't exist in Vista or XP, Vista and XP machines wouldn't know what to do with them. Now if you do have an environment that includes those older clients, what you need to do is this: you can download something called Client Side Extensions from Microsoft's Web site--it's a free download. And you would roll these out to all of your Vista or Windows XP machines, and you can use Group Policy to do that if you use a regular Group Policy.
Once the Client Side Extensions are installed, basically they do exactly that. They extend the clients so that they can process the Group Policy Preferences. I think you're really going to enjoy these. Group Policy Preferences are really cool, and they're going to give you lots of functionality that we've been waiting for, for years.

8.10.2 Configuring Group Policy Preferences

Configuring Group Policy Preferences
0:00-0:07
In this video, we're going to take a look at Group Policy Preferences. Of course, we need to get into Group Policy.

Group Policy Management Console
0:08-0:37
We're going to go into Group Policy Management Console. Preferences are only set in Active Directory, so you will not see them in Local Group Policies, just in the Group Policies that you create in Active Directory. I'm going to get Group Policy Management Console out of the Tools menu, but you could get it from the Start menu as well. Since we're just really interested in looking at the policy, I'm just going to edit the Default Domain Policy so that we can take a look.
Group Policy came in with Windows Server 2000 and Active Directory.

Group Policy Preferences vs. Logon Scripts
0:38-1:12
It gives us the ability to centrally manage settings on the workstation. Before that, the only thing administrators could use were logon scripts, and even when Group Policy came in, it still wasn't able to do everything that we could do with logon scripts. Even, let's say, in the Server 2003, 2008--there were certain settings that could only be done with a logon script, and you'll still see logon scripts in the environments if any of these things needed to be done.

Group Policy Preferences
1:13-1:19
So with Preferences, Microsoft wanted to go in and address some of those features that were lacking.

Computer Configuration Preferences
1:20-1:33
You can see I've got Preferences on the computer side, and I've also got Preferences on the user side. There are a lot of settings that overlap, but some that are specific to the user or specific to the computer. I can come in.

Windows Settings
1:34-2:00
You can see under Windows Settings I can go through and do quite a number of things. I can set up environmental variables for applications. I could actually copy a file down to the local workstation. I could create a folder on that computer. I could adjust the Ini files. I could go through and add a setting to the registry. I could create a share. I could put a shortcut anywhere on the computer, probably on the desktop.

Control Panel Settings
2:01-2:22
I've also got Control Panel Settings, so I can add ODBC data sources, add devices, different folder options, create local users and groups, set up network options, power options, install a printer-- this would be a local printer-- schedule tasks, adjust services, so quite a lot of functionality in there.

User Configuration Preferences
2:23-3:04
It gets even cooler on the user side, where I can go through and I can map drives, which is a big one that had to be done inside of logon scripts. You can adjust applications. You see some of these are the same as we see up in the computer side, and then down here I've got internet settings that I can set up where I can do settings that affect different versions of Internet Explorer. Again, I can create local users and groups, network options, power options, printers. I can even adjust their Start menu. So, a lot of flexibility coming in with Group Policy Preferences.

Client Side Extensions
3:05-3:39
The only sort of catch to this is it did come in with Server 2008 R2 Windows 7. If you have older clients, the preferences are not going to affect them. Even Vista--it won't affect that computer. What you can do is go out to Microsoft and download something called Client Side Extensions, and it's just a little add-in software that would be installed on the workstation. You can roll it out via Group Policy. Once the Client Side Extensions are installed on those XP and Vista machines, then the preferences will affect them, and you can use them throughout your environment.

Summary
3:40-3:52
So, Group Policy Preferences-- really coming in and giving us a lot of additional functionality in adjusting the user environment, so that we can finally get rid of those pesky logon scripts. And that's how you configure Group Policy Preferences.

8.10.3 Group Policy Preference Facts

Group Policy preferences, introduced with Windows Server 2008 R2, allow you to configure, deploy, and manage operating system and application settings. Group Policy preferences are processed only on Windows 7 and Windows 8 client computers.

To apply Group Policy preferences to pre-Windows 7 clients, download and install client-side extensions (CSEs). You can use Group Policy to rollout the CSEs to pre-Windows 7 client computers.

The following table provides an explanation of available Group Policy preferences:

Preference	Description
Drive maps	Manages network drive mappings without writing logon scripts.
Environment	Manages user and system environment variables or updates the environment path.
Files Folders	Manages files or folders, such as copying configuration files to users' profile folders, or regularly cleaning up temporary folders.
Ini Files	Modifies and updates individual properties within a .ini file.
Network shares	Manages network shares on multiple, targeted computers. Additionally, it configures user limits and prevents users from seeing subfolders for which they lack permission to access.
Registry	Manages registry entries without the need to write scripts.
Shortcuts	Manages several types of shortcuts on multiple, targeted users and computers.
Devices	Enables or disables devices based on a device class identifier.
Folder options	Configures folder options and file extension associations.
Internet settings	Configures Internet Explorer options for Windows Internet Explorer.

Local users and groups	Manages local users and groups.
Network connections	Configures VPN and dial-up connections.
Power options	Configures power options and power schemes for computers.
Printers	Manages shared printers, TCP/IP printers, and local printers.
Regional options	Configures the user locale, including number, currency, time, and date formats.
Scheduled tasks	Manages scheduled tasks on targeted users and computers.
Services	Configures services to: • Run automatically • Start if required • Disable and stop if necessary
Start menu	Configures Start menu options for users.

9.1 IPv4 Addressing

As you study this section, answer the following questions:

- What is the format of an IPv4 address?
- What is the purpose of a subnet mask?
- What is the default address class of the IP address 132.11.166.5?
- What is the decimal equivalent of the IP address 132.11.166.5?

After finishing this section, you should be able to complete the following tasks:

- Configure static IPv4 settings.
- Convert an IPv4 address from binary to its decimal equivalent.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  Configure Basic Network Settings
  - Configure IPv4 Settings

This section covers the following 70-410 exam objective:

- 401 Configure IPv4 and IPv6 addressing.
  This objective may include but is not limited to:
  Configure IP address options

9.1.1 IPv4 Basics

IPv4 Basics
0:00-0:07
We're going to talk about TCP/IP, specifically IPv4. TCP/IP is the protocol of the Internet.
Protocol
0:08-0:51
Protocol is the language that computers speak, so if I'm teaching this video in English and you speak English, great. We'll have communication. We're using the same protocol, but if I were to teach it in French and you don't speak French, not a whole lot of communication going on.
Since the beginning of the Internet, they've used TCP/IP, and specifically IPv4, and we're going to go over the basics of IPv4 and get a handle on how these computers use it to communicate with each other. In TCP/IP, any device that has an IP address we call a host. That's the word I'm going to try to use.
Let's take a look at the white board and see the major rules of TCP/IP that are never violated. The first rule of TCP/IP-- and these three rules, you can completely take them to the bank, they're never, ever, ever violated-- is that every host on a network must have a unique IP address.
Rule 1: Unique IP Address
0:52-1:14
Those people already know that when they come into it, if you had experience with TCP/IP before. I think of them like a cellphone number. If I call your cellphone number, only your cellphone should ring. If another person's cellphone rings at the same time, now we have a problem.
Rule 2: Same Network ID
1:15-1:45
Second rule, all the hosts on the same network must have the same network ID. We can think of the network ID as the identifier for that particular network. In the US, we use area codes. For example, in Rhode Island, the area code is 401. If the number starts with 401, it's a Rhode Island number. If it doesn't, it's not a Rhode Island number. Same thing with the network ID. All the hosts on the same network must have the same network ID, because if they didn't have the same network ID, they wouldn't be on the same network. We'll dig into that a little bit more in a few minutes.
Rule 3: Direct Communication Within Own Network
1:46-2:15
The third one is really the most powerful rule, again, never ever violated, but we're going to have to take a while to really understand it, but hosts can only communicate directly with other hosts on their own network. You might say, well, I know my computer routinely communicates with computers that are not on my network. Sure, absolutely. But they do that not directly, but through other devices called routers. We're going to go through these rules and get a little bit better understanding of them.
TCP/IPv4 Addresses: 4 Octets and 32 bits
2:16-2:35
TCP/IPv4 addresses, their IPv4 addresses, are made up of four octets. These would be eight-digit binary numbers. They're written in decimal because most people are not very skilled or don't have a lot of bonding with binary, but it gives us a total of 32 bits.
That will come back again later, but it's something you should be aware of.
Example TCP/IPv4 Address
2:36-3:23
Let's go ahead and get an example of a TCP/IPv4 address. I'm just going to pick a very simple one. I don't know why this is my favorite IP address, but it is.
We're going to do an IP address of 192.168.1.10. We look at that IP address and we say that represents some particular host on a network with an area code or phone number. I had (401) 555-1212. Most people will say, okay, well, the area code is 401, because that's the first three digits. Very easy with a phone number.
IP addresses are not so easy. The numbers that represent the network ID are sort of up for grabs. By looking at an IP address, there's really no way to tell what network ID this is on--which part represents the network and which part represents the individual host on that network.

Subnet Mask

3:24-4:15

When I teach the long classes, sometimes I'll tell my students, there's only one trick question I'll ever ask you, and that's what network is this IP address on, and the only correct answer is, "I don't know, Shad, give me the subnet mask." Because you can't identify the network ID unless you have a subnet mask.

The subnet mask is another 32-bit binary number, so we'll give it a subnet mask. Its only job is to tell the computer which bits in the IP address represent the network ID and which bits in the IP address represent that particular host. We're concerned with the network ID because we know we can only speak directly with other computers on the same network.

When this computer wants to contact another computer, it has to know. Are we on the same network? We're going to talk directly. Or are we on a different network? And then I have to go through a router or some kind of device that's going to translate this message over to the correct network.

Right now, we're going to keep it very simple.

ANDing Process

4:16-5:52

The rule at this stage of the game is this: Any number in the IP address that's in the same position as a 255 in the subnet mask, so we can see 192--that's my first octet. In my subnet mask, I've got a 255 in the first octet, that is a part of the network ID.

If I were trying to write my network ID here underneath the line, I would write 192 as definitely being part of the network ID. When you're doing this, I would suggest that you always go through this process. It's actually a process called ANDing. Write your IP address; try to line up the periods as I have. I haven't done a wonderful job, but I sort of got them lined up, and go through and find the network ID.

I usually find that no matter how simple it seems, if I go through this way, I get the right answer. As soon as I start feeling clever and I shortcut, then I have problems. Here, a 168 in the second octet, same position as a 255, so that's going to be part of my network ID.

One, third octet, same position as a 255, that's going to be part of my network ID. Now we get to the fourth octet. This ten is in the same position as a zero. Anything that's in a same position as a zero is that host on the network. It's not part of the network ID, but we write a zero to symbolize that fourth octet.

There's no such thing as a 192.168.1 network. This is the 192.168.1.0 network. Any IP address that we would look at in this network with this particular subnet mask should come out to having that network ID. If it doesn't, it's not on my network. Again, we're only going to communicate directly if we're on the same network.

Examples of Computers Communicating Directly

5:53-5:57

Let's take a look at an example of some computers that would communicate directly.

Local

5:58-8:12

I have a sending computer over here. I have given it 192.168.1.4, with a subnet mask of 255.255.255.0. I'm sticking with the same one I just used. Again, we're going to go through the same process. 192 same position as a 255, that's part of my network ID. 168 same position as 255, that's part of my network ID. 1 same position as a 255, part of my network ID.

In the fourth octet, I have a zero in the subnet mask, so that is not part of my network ID. This particular computer--we can even call it computer A--with 192.168.1.4 as the IP address, 255.255.255.0 as the subnet mask-- is on the 192.168.1.0 network.

Now, let's suppose that computer A wants to talk to computer B. The first thing it has to find out is, is computer B on my network? If it is, we're going to talk directly. If it's not, then we've got to do something else. Computer A does not know what computer B's subnet mask is, but the assumption is if computer B is on the same network as computer A, then it would be using the same subnet mask.

You always use the sender's subnet mask to evaluate the receiver, so because A has 255.255.255. Do a better job of lining up your periods than I did. All right. We're going to use that same subnet mask over here and go through the same process. So 192--looks like it's part of the network ID--168.1.0. If these numbers are equal, these computers are on the same network.

In real life, that tells the computer, "Well, we must be connected to some type of a device." All right. We can draw a little switch in here and this is connected. What computer A is going to do, is going to go through a process to find out the MAC address of B and just send that packet out.

We're on the same network. It's going straight to network B. I don't have to do anything fancy. When they come out the same like that, it's great. If we were troubleshooting a situation like this, if these two computers couldn't talk to each other, we would know that the problem is in our network.

Maybe it's with the network card on A or the cable between there, or maybe it's a problem with the switch or the cable or the network card on B. They should be talking directly to each other. When they're on the same network, we call them local.

Remote

8:13-9:08

My first question to myself is, are these computers local or remote?

Remote would be if we're not on the same network. Let's take a look at an example like that. We've got it set up here. I've got a nice computer here. We'll give it a subnet mask of 255.255. To make it interesting, we'll do 0.0. Then, after we do our basic ANDing, we come up with a network ID of 192.168. In the third octet, there's a zero, so that's going to be 0.0.

Now, most people would say, "Hey, Shad, I can tell by looking that that other computer is not on the same network." I never, ever guess like that. I know you can probably tell right away this is going to be different, but I still do the math. I find that the more methodical I am, the happier I am when I'm troubleshooting.

So again, we use the sender's subnet mask to evaluate the recipient. Of course, we're going to get a completely different network ID, 214.32.0.0. These are not equal, therefore they are remote.

Routers

9:09-12:16

If the computers are remote, they're on different networks. We have to have some sort of a device to connect the networks together.

The devices that connect networks together are called routers. I always draw a router circle with an X in it. A router is basically any device that's connected to two or more different networks and can pass information between them. Knowing it, finding out that these numbers are different, tells us there's at least one router in between.

In order for the computer to be able to talk with other computers in different networks with remote computers, it has to have a Default Gateway. The Default Gateway is the IP address of the router in that computer's network, because it can never violate that third rule. I can only speak directly with other hosts on the same network as myself.

Assuming, let's say, the default gateway given here is 192.168.1.1. I would then do the math again, because that's what the computer's going to do. It's going to take its subnet mask, put it up against the default gateway, do the math.

I don't have too much room here, but I'm going to find out that, again, it's going to be this: 192.168.0.0. Any computer that's remote from me, my computer's going to say, "Send it to the gateway." I always think of it the way children say, "Mom"; anything they don't know how to do, "Mom ..." That, to me is the router's default gateway.

Now, this router has got to be connected to two or more different networks and be able to pass information between them. This router here--we'll call it router 1--is going to get the packet. It says, wait a minute. I'm connected to 192.168.0.0 network, right here on this side. On this side, I'm connected to a different network.

Let's say the network here is 192.168.2.0 and this is .1. This router here is .2-- router 1 gets the packet. On this side, I'm 192.168.0.0. On this side, I'm 192.168.2.0. Neither of those is 214.32.0.0, which is where this is headed. Let me give the packet to my Default Gateway, some other router that's local to me that's further along in the chain.

Let's say router 2 gets it and says, wait a minute, 214.32.0.0--not my network, but I'm connected to some other network with another router. We can go ahead and put another router in here. I'm going to pass the packet along to them, until finally we get to a router that says, 214.32.0.0. That's my network. I can talk directly to whoever this packet is intended for and give it directly to them.

That's exactly how the Internet works. I get no end of pleasure thinking of it sort of like the Pony Express from the Old West, or a baton relay race. All these computers talking locally, each one is local to each other, and yet we cross vast distances.

Sometimes, there might be--if this is my home network, there might be 10 feet between me and my router. Up on the backbone of the Internet, there could be thousands of miles between router 2 and router 3. But that's essentially how it's going to go.

Our first question: are we local or are we remote? If our network IDs come up exactly the same as they did here, then we're local. We're going through a switch. We're talking directly. If our network IDs come up different, then we're remote. We're going to go through a series of routers.

IPconfig /all Command

12:17-12:29

Our very best commands, first of all, IPconfig. I'd like to add the /all, because that will give me my default gateway and my DNS.

That will let me evaluate what network ID my computer is on.

Ping Command

12:30-12:40

Then, I normally will use a ping. Then, you can put the name of the computer that you're sending it to-- should give you the IP address of the other computer. You can go ahead and do the math on that and find out if you're local or remote.

Tracert Command

12:41-13:16

Then if you think there's a problem, you find out it's remote. You could also use the command, tracert. You say it trace route, but we type it tracert and then, again, I could either put the name, or I could put the address. Remember, we had 214.42.80.30, whatever it is, either way. What that does essentially is, it does the same as ping, but it says every router that you have to go through between your computer and the remote computer, have that router send back its name and IP address, so we can trace the route that this is taking through the Internet.

Summary

13:17-14:12

Those are IPv4 basics. Again, the most powerful ones are our rules up here. Every host on a network must have a unique IP address. All the hosts on the same network must have the same network ID. We've talked about how to find the network ID. Then hosts can only communicate directly with other hosts on their own network--that's our local or remote.

Okay, so those are the basics of TCP/IP, version 4. We went through the three main rules that you need to make sure you remember. We did our basic ANDing. Whenever you come to an IPv4 problem, always ask, "What network is my computer on? What network is the other computer on?"

If we're local, we're looking for a problem inside the network. If not, what network am I on? What network is the default gateway on? If we're local, great. I'm going to see, make sure where the problem is. If I can talk to my default gateway, maybe I'll use a tracert to see where in the chain of routers is breaking down. Those are the basics of IPv4.

9.1.2 Binary Numbers

Binary Numbers
0:00-0:30
We're going to talk about the binary number system. When we write out IPv4 addresses, we write them out in decimal numbers. The computers actually process them as binary numbers, so it's important to be able to convert from decimal to binary or binary to decimal. Either way. Then once you know how to do that, you can go through and simulate the calculations the computer is going to do, so that you can troubleshoot if anything is going wrong with an IPv4 address.

Decimal Numbers
0:31-2:10
Let's take a look at a decimal number first, and then we'll look at a binary number.
What I found is that it's easier to understand binary if you first take a look at what you already know with the decimal number system. Because all number systems work the same way, they just have a different base. The decimal number system that you use when you go out grocery shopping or you go get something to eat, that's the decimal number system that's base 10.
I've put up a decimal number here, 4389, and hopefully everybody would agree that this is 4,389. Most people then could say, OK. That is really the same as 4 x 1,000 + 3 x 100 + 8 x 10 + 9 x 1. You're probably saying, "Shad, that is so 2nd grade man, get to the point." All right.
We also then agree that this is really the same as saying 4 x 10 cubed + 3 x 10 squared (10 Cubed is 10 x 10 x 10. The 10 x 10 is 100; x 10 is 1,000, 10 x 10 is 100) + 8 x 10 to the 1 (anything to the 1 is itself) + 9 x 10 to the 0. That's the one most people haven't heard of. Anything to the 0 is really the number 1.
All I'm doing in this second line is really repeating what I did in the first line.

Decimal Number Table
2:11-3:20
Once you say, OK, I've wrapped my mind around that, we could actually draw out a little table to describe a decimal number system and say, OK. The first slot is 10 to the 0, and then we have 10 to the 1, and then we have 10 squared, and then we have 10 cubed. This is really the number 1. This is really 10. This is 100. This is 1,000. I've got four of these, three of these, eight of those, nine of those.
Any number system can be described with a chart like this, using the appropriate base. That's the first concept of number systems.
The second concept is in any particular column, I can have one digit, and I can go from 0 up to 1 less than the base. Since we use the decimal number system or the base 10, I can go 0 through 9, and hopefully everybody understands that when I get back up to 10, I've got 4389, I add 1. I'm going to roll back around to 0 and increment the next column. I should get 4390. That's the decimal number system.

Binary Numbers
3:21-3:30
Really, if you have a good grasp of that, all you really have to know about the binary number system is it's a base two. It follows exactly the same principles.

Binary Number Table
3:31-4:21
So let's look at some binary numbers.
IPv4 addresses are made up of octets, which are eight-digit binary numbers. When we put together a grid to describe the binary number system, we want to make sure that we have eight positions.
OK. We're going to start with 2^0 right here, and go all the way to the left through 2^7, because we started at 0, and we go to the 7. That means we've got eight places that we can go through. Anything to the 0 is 1. 2^1 is 2 times 2, is 4, times 8, 16, 32, 128, 64, et cetera.
In any of these columns, we can go up to one less than the base. But here, because we've got base 2, the only thing you're ever going to see in a binary number is either a 1 or a 0, because if I get to 2, I'm going to roll back to 0 and increment the next column.

Example: Convert Binary to Decimal

4:22-5:33

I can go through in these numbers, and let's just make up a binary number. It doesn't matter to me what it is. This is the same as saying 1 x 128 + 0 x 64. I could write the whole thing out, but we know that anything times 0 is 0. I don't really need to deal with those columns.

The easier way to process this number is to just add up all the spots where we have 1s. Some people like to think of 1 as on, 0 as off. I tend to be a little bit more mathematical, but whatever works for you is always the right answer. I can go in here and say, OK. I've got one 128, I've got a 32, I've got a 4, and I've got a 2. That's 10, 4 ,and 2 is 6, so I've got 16 here. That gives me another 6 with a 1. So this binary number, 10100110, is the same as the decimal number 166. I could come up with any combination and take that binary number and make it into a decimal number.

Second Example: Convert Binary to Decimal

5:34-6:48

Just to make sure, we'll do that one more time. Get our grid up here. Anytime you have to deal with binary, just put your grid down. The hardest thing to remember is that your 2 to the 0 is 1. 2 to the 1 is 2, and then you just multiple by 2 going left. Let's pick 111011010, whatever that is, that give us one 128, a 64, no 32s, one 16, one 8, no 4s. I've got one 2 and no 1, so I'm going to have 10 there, plus another 10 is 20. It looks like 28 to me. Ironically, I'm pretty good with binary, terrible with this kind of math. All right. So that looks like 11. This is 218 in decimal. We can go through and get any number that we want.

If all of these numbers were 1s, that's where we get our 255 from. If all of these four numbers were 0s, that would still be a 0. That's why numbers in an IP address range from 0 to 255, because that's our minimum and our maximum decimal numbers given an eight-digit binary number.

Example: Convert Decimal to Binary

6:49-8:36

When you're working with IP addresses, you may also have to convert from decimal to binary. Not always binary to decimal. We're going to take a look at how to do that as well. Again, anything to do with binary is going to start with our grid, and we'll pick any number between 0 and 255. I'll do 172. All right.

Always start on the far left and what you're asking yourself is, is this number 172 greater than or equal to that position of 128? In our case, it is. If the answer to that is yes, greater than or equal to, I'm going to put a 1. There's one 128 in 172. Now I'll subtract my 128, and that should give me 44 left over.

Is 44 greater than or equal to 64? No. If the answer is no, I'm going to put a 0.

Is 44 greater than or equal to 32? Absolutely. I've got one 32 and 44. Now I'm going to take my 44, subtract 32 and I should get 12 left over.

12 greater than or equal to 16? Nope.

12 greater than or equal to 8? Absolutely. I've got one 8 in there. I subtract my 8 from 12. I've got a 4 left over.

4 greater than or equal to 4? Absolutely. I still subtract my 4. Now, I'm working with the number 0.

Is 0 greater than or equal to 2? Nope. Is it greater than or equal to 1? Nope.

So the number 172 in decimal is going to be 10101100 in binary.

Second Example: Convert Decimal to Binary

8:37-9:56

Just to make sure, we'll do that one more time. Let me pick something a little smaller, something like 83, so we can see what happens. You always want to work with all eight columns because, again, IP addresses are eight-digit binary numbers, so even if it starts out with some 0s, that's OK.

Is 83 greater than or equal to 128? Nope. Put a 0 there.

Greater than or equal to 64? Absolutely. I've got one 64. Do my subtraction here. That's going to give me 19.

19 greater than or equal to 32? No way.

19 greater than or equal to 16? Absolutely. We put a 1. OK. I subtract my 16 from 19. I've got 3 left over.

3 greater than or equal to 8? Nope.

Greater than or equal to 4. Nope.

Greater than or equal to 2. Absolutely. I've got one 2, 3 minus 2 gives me 1.

One is absolutely greater than or equal to 1.

My decimal number 83 is the binary number 01010011.

That's how you convert from decimal to binary or binary to decimal.

All right.

Summary
9:57-10:17
We've gone through the binary number system. First, we looked at the decimal number system just to remind you what you already know. Then, we looked at the binary number system, going both from binary to decimal and decimal to binary. Just make sure you understand how to write down that grid. Make sure you get all eight columns and you'll be good to go with IPv4 addresses.

9.1.3 IPv4 Classes

IPv4 Classes
0:00-0:07
All right. In this lesson, we're going to talk about the IPv4 Default Classes and NAT--network address translation.

IPv4 to IPv6
0:08-1:41
Originally, when the internet started, it was all public IP addresses. I remember being a very young instructor around 1997, 1998, and all textbooks would say, we're running out of IP addresses, IPv6 is coming out. IPv6 is coming. It's coming. It's coming next year.
Well, it's more than a decade since then, and we're still not quite there. They had to do something in the interim to stop the depletion of IPv4 addresses so that they could get the world ready for IPv6.
One thing that a lot of people don't understand is for computers to use IPv6, it has to support IPv6 from the sender all the way to the receiver. That means the sender's network card, the sender's switches, routers, all the routers on the internet; same thing at the recipient's side.
Back in the late 90s, people had just spent sometimes millions and billions of dollars putting in equipment that supported IPv4. We don't really want to go back to these companies and say, "Well, glad you spent $34 billion implementing a network, but we need you to junk it all and get things that are compatible with IPv6."
At that time it was InterNIC, now the organization that runs the internet is IANA. These guys came up with an idea to put off the depletion of IPv4 addresses so they could get a window of time to get the world ready for IPv6, and it worked really great.
Let's take a look at the IPv4 classes of the public addresses. We'll look at the private addresses and we'll see how they got this system to give them a little bit of time.

Classes
1:42-2:34
I've drawn just a really simple network and we know that we have our sender here, computer A. It's going to go out to some server on the internet, could be Yahoo.com, it could be MSN. It doesn't matter. It's going to go through and say, "Am I local or am I remote?" Well, definitely, they're not on the same network; they're remote, so it's going to send it to its default gateway, the router, the address of the router on its network. That router is going to send it to another router, to another router, until it finally arrives at the destination. That's how the internet works.
Originally, there weren't a lot of hosts on the internet. It started out as a Department of Defense project and, for whatever reason, they said, "Let's divide up the IP addresses that will be used on the internet into classes."
There are three classes that are used for hosts: Class A, B, and C. There are a couple other classes that are used for other things, but this is really all we're concerned with.

Class A
2:35-3:05
They said, "Let's make Class A, be any IP address that starts with a 0. Mathematically, I could have a 0 and then seven more 0s, which will give me the number 0; or I could have a 0 and then seven 1s, which would give me the number 127. But they said, let's not use 0, let's use 1.
And 127, they took out for testing, so Class A is any address that starts with the numbers 1 through 126. If that very first octet is one of those numbers, it's a Class A address.

Class B
3:06-3:19
Class B, they said, let's let that be anything that starts with a 10 in binary, which would take me up through 10, and then six 0s or 10 and six 1s, which are the numbers 128 through 191.

Class C
3:20-3:30
They said, "Let's have Class C be anything that starts with 110." Well I could add five 0s after it or five 1s, which would give me the numbers 192 to 223.

Ranges
3:31-4:06
You don't have to know the binary. I actually have a lot more bonding with this system where they started with 0 and they keep adding 1, and so on. You do need to know their decimal ranges. Whether you're like me and you work it out in binary this way, or if you just memorize the ranges-- either way, you should know them. If you don't love memorizing, I would say at the very least, memorize the range with Class B, as if it's below that, it's a Class A; and if it's above it, it's Class C.
There's actually a D and E beyond that, but that's reserved for multicasting and experimentation.
I've got my classes here and here are the numbers; this column here are the number ranges again.

Default Subnet Mask
4:07-5:19
They assigned each of the class a default subnet mask.
If it's a Class A, it starts with the numbers 1 through 126. Then they gave it a default subnet mask of a /8.
Class B, meaning the very first octet is 128 through 191-- they gave it a default subnet mask of a /16.
Then Class C, any one where the first octet is 192 through 223-- the default subnet mask is a 224.
The default subnet masks--these are the numbers of 1s in the subnet mask-- if I'm using eight 1s for the network, that means out of 32, I have 24 leftover for hosts. It's 16 million clients all on one network. These are huge networks, which you would then use subnetting to break down into smaller ones. They were given out mostly to ISPs and big organizations like that.
Class B, if I'm using 16, then I've got 16 left over for hosts. That gives me 65,534 clients, but again on one network-- still very big, but usable.
Class C is if I'm using the first three octets for my network. I've only got eight left over for hosts, and that's our familiar 254 clients all on one network.

The Problem and Solution
5:20-6:52
At the time, when the internet first started, it made things very difficult. Now what do you do if you're a small company and you have, let's say, 1,000 computers, and you have 1,000 spread out over four buildings? Do you go by four Class C addresses, and then you're roughly there? You've got the right networks, but now you've got about four networks from InterNIC. Or do you buy one Class B and then just waste 64,000 addresses? It was causing a lot of waste of IP addresses, and that's why we were depleting the IPv4.
The powers that be on the internet said, "Well, here's what we're going to do. We're going to pull some networks off the internet and then grab ranges from each of the classes."
So in Class A, anything that starts with a 10, they said, "We're going to pull that off the internet. We're going to guarantee you that no matter what website you go to, you are never going to be trying to go somewhere on the internet, and it's going to start with a 10, and it's going to look like it's in your basement."
In Class B, they took 172.16, all the way up through 172.31. There were a bunch of Class B networks that came off the internet.
They were really generous in Class C. They did anything that starts with 192.168, which is essentially 255 Class C networks that they pulled off the internet. They said, "We're going to guarantee that these will not be used on the internet, which means you can use them privately in your companies or in your home and you will never have to worry about it looking like Yahoo is in your basement."
They did pull a fourth range off, a 169.254 network from Class B. That's reserved for APIPA, which we'll talk about in our DHCP Chapter.

Private IP Addresses
6:53-7:26
Now with these private IP addresses, I can go ahead and use them on my network. Most home networks probably have something that starts with a 192.168. Let's say you're on 10/24 network, so maybe this computer here is .10, and maybe I even have another computer here, and it picked up 192.168.1.11, and then I added a laptop, and that got .12.
It's great that they pulled it off the internet, which allows us to use these networks IDs in our private homes.

Network Address Translation (NAT)

7:27-10:28

But then comes the problem, well, hey, if I have the 192.168.10 network and my cousin has the 192.168.10 network and all his friends have the same network and so on, how does the computer know which one of us to give the answer back to? Because again, if we go back to our rules of TCIP, every host has to have a unique ID, they all have to have the same network ID. How is it going to get back if everybody's using that network ID?

The answer to that is something called NAT. NAT stands for Network Address Translation. When you have these private addresses, you have to have a NAT router. A regular router is any device that's connected to two or more networks. Here, my router is connected to two or more networks, and can pass information between them. They really are just like a traffic cop; if it comes in from network A, I hand it over to B; if it's headed for B, if it comes in from B, headed over to A, I just hand it back to A.

A NAT router works differently. In a NAT router, one of the interfaces is on a private network. Of course, that would be this one. Let's say my little home router here picked up .1, so it's 192.168.1.1. The other interface is on a public network. Well, really the only public network is the internet, so this is out on the internet. Let's just say, for the sake of argument, I got 63.120.13.4. It doesn't matter what it is. As the computers in my house make requests, the NAT router takes the request from the private interface and repackages it, so everything that's going out seems to come from this one public address.

Let's say .10 asks for Yahoo.com, and .12 wanted MSN, and .11 wanted NBC. All those requests would go out through this NAT router, but all three of those requests would come from this one public address. Because it's a public address, when it hits that server, whether it's Yahoo, MSN, NBC, that server knows how to get back to that public interface, and then the NAT router keeps a little list. Well, you know, Yahoo is coming back, let me give it to .10, NBC is coming back, let me give it to .11, and so on and so forth. Because that net router repackages it, again, Network Address Translation translates the private IP addresses into public IP addresses; that allows you to have a theoretically unlimited number of private addresses behind one public one.

Well now, if I'm a company and I have 1,000 computers, I can get one internet address from my ISP. That's where this is coming from, and have my 1,000 computers behind that, or I can go out and get three addresses from my ISP, and have my three sites with 50 or 500 computers, or whatever I've got at each one. It's no big deal, because this NAT will take care of making sure that things go out on the public address and come in to the public address. Now we can put off the depletion of IPv4 addresses and get ourselves another decade or 15 years before we have to go to IPv6.

Summary

10:29-10:57

That's a little bit about the IPv4 default classes out on the internet: Class A, B, and C. We should know the decimal numbers, be able to look at the first octet of an IP address, and place which class of network it's in. You should know the default subnet mask for each of those networks. You do not need to memorize how many hosts; that was just for fun. Have an an understanding of Network Address Translation, how the NAT routers have one private interface, one public interface, and everything that goes through them gets repackaged with the IP address of the public interface.

9.1.4 Configuring IPv4

Configuring IPv4
0:00-0:40
In this video, we're going to take a look at configuring TCP/IPv4. There's a number of ways to get in here. One way would be to go to the Local Server. I can click on Ethernet and that will take me right to my Network Adaptors. If for some reason you don't like that method, you can come all the way down in the bottom right hand corner to the icon that symbolizes the network, right click, and go to the Network and Sharing Center. From the Network and Sharing Center, I will Change adaptor settings.
Somehow you've got to get to where you can see the adaptor.

Properties of TCP/IPv4
0:41-1:10
Once I'm at my adaptor, I will right click and choose Properties. I then need to go into the Properties of TCP/IPv4. This particular computer has a static address. That means I've typed in the address in the subnet mask. It doesn't have a default gateway, which means it can't get out to the internet. It's got to have a default gateway if it's going to get out of its own network. It also has a preferred DNS server.
If I wanted to use DHCP--now watch this--right now, I only have a General tab, I would obtain an IP address automatically, and that puts up my Alternate Configuration tab.

Alternate Configuration Tab
1:11-2:25
This radio button is what causes the computer to pick up an APIPA address if the DHCP server is not available. These are addresses that start with 169.254.
If I have a computer that needs to be functional both in environments that have DHCP and on networks that have static IP addresses, so maybe this is a laptop that needs a static IP address at work, but uses DHCP at home. In that case, I can be clever. I keep it on Obtain an IP address automatically here, which says I'm going to try for DHCP, but then for my alternate configuration I would set up a static IP address that would be valid at work. When the user goes home, they'll pick up an IP address via DHCP. When they come to work after DHCP doesn't respond, it will fail over to the correct static address.
I also have some advanced settings that I can set up for TCP/IPv4.

Advanced Settings
2:26-2:52
If I did need to have multiple IP addresses or multiple default gateways, I could add them in here. It's not necessarily advised, but we call that a multihomed computer. It's any computer with more than one IP address. Maybe this is a network adaptor that is connected to a couple of different networks. There are situations like that.

Multiple Default Gateways
2:53-3:10
Multiple default gateways means that I have multiple ways to access the internet. In that case, if you're going to add a backup router or backup default gateway, make sure that the metric of the backup router is higher than the primary one. I have some advanced DNS settings.

Advanced DNS Settings
3:11-4:20
I can add in different DNS servers to use. By default, the computer will pin the primary in connection specific suffixes to see if I go through and I ping a single label name to see if it can get a good, Fully Qualified Domain Name.
I could set up a list of suffixes to try. I can set a DNS suffix for this connection--register the connection's address in DNS. If I didn't want this computer to register with DNS for whatever reason, I could uncheck this and it would not register with DNS. If I put in a connection specific suffix, let's say this computer is a member of northsim, but I want this particular adaptor to register itself with eastsim, I could put in a suffix and tell it to use the suffix in DNS registration. This is generally actually done through DHCP and not manually, but I can do it manually in this tab.
The WINS tab would be used if WINS is available in my network.

WINS Tab
4:21-5:00
WINS is an automatic way of resolving NetBIOS names. NetBIOS names have not really been in use since NT4.0, so hopefully this is not something you ever have to get involved in. By default, you can actually use NetBIOS if DNS doesn't work out. In my entire career--more than 18 years--I haven't had to come in here and adjust this, but if you do need to, it's available for you.
That's how we set up TCP/IPv4. For static IP address, you just type it in. I hit okay and I'm good to go.

9.1.5 IPv4 Facts

IP addresses, in conjunction with routers, are responsible for sorting and delivering packets to and from clients on a network. Each packet contains the IP address of both the sender and the recipient. Routers use the IP address to send the packets to the specified destination. IPv4 addresses allow hosts to participate on IPv4 based networks.

Keep in mind the following rules for using IPv4:

- Each host must have a unique IPv4 address.
- Each host on the same logical network must have the same network address.
- Hosts can communicate directly only with other hosts on the same logical network.

Be aware of the following IPv4 concepts.

Concept	Description
Host	A *host* is a computer on a network
IP Address	The *IP address* is a number assigned to identify hosts and other components on a network.
Network address	The *network address* is the portion of the IP address that identifies a specific network. The remaining portion of the IP address identifies the host or other component on the network.
Subnet mask	A *subnet mask* identifies the portion of the IP address that defines the network address and the portion of the IP address that defines the specific host.
Address Class	IPv4 addresses are divided into *classes*. The *address class* identifies the range of IPv4 addresses and a default subnet mask used for the range.
Default subnet mask	A *default subnet mask* is assigned to classes A - C as follows: • **255.0.0.0** is the default subnet mask for class A networks. • **255.255.0.0** is the default subnet mask for class B networks. • **255.255.255.0** is the default subnet mask for class C networks.
Broadcast address	The *broadcast address* is the last address in the address range and is used to send messages to all hosts on the network.

Default gateway	The *default gateway* is a device that performs the act of routing and enables a host to communicate with other hosts on other networks through the process of routing. • A default gateway address must be configured on each host to allow inter-network communication. Without the default gateway, hosts will be able to communicate only with devices within the same subnet. • The default gateway address must be on the same subnet as the host computer. Routers have multiple network interface cards attached to multiple networks. When configuring the default gateway, choose the address on the local subnet.

IP Address Structure

An IPv4 address is a 32-bit binary number represented as four octets (four 8-bit numbers).

- Each octet is separated by a period. IPv4 addresses can be represented in one of two ways:
  Decimal (for example 131.107.2.200). In decimal notation, each octet must be between 0 and 255.
  Binary (for example 10000011.01101011.00000010.11001000). In binary notation, each octet is an 8-character number.
- To convert from binary to decimal and vice versa, memorize the decimal equivalent of the following binary numbers:

1000 0000	0100 0000	0010 0000	0001 0000	0000 1000	0000 0100	0000 0010	0000 0001
128	64	32	16	8	4	2	1

To convert from binary, take each bit position with a 1 value and add the decimal values for that bit together. For example, the decimal equivalent of 10010101 is: 128 + 16 + 4 + 1 = 149

Subnet Mask

The *subnet mask* is a 32-bit number that identifies the network portion of the of each IPv4 address.

- In binary form, the subnet mask is always a series of 1's followed by a series of 0's (1's and 0's are never mixed in sequence in the mask).
- A simple mask might be **255.255.255.0**.

The following table shows the default address class for each IPv4 address range.

Class	Address Range	First Octet Range	Default Subnet Mask	CIDR Notation
A	1.0.0.0 to 126.255.255.255	1-126 (00000001--01111110 binary)	255.0.0.0	/8
B	128.0.0.0 to 191.255.255.255	128-191 (10000000--10111111 binary)	255.255.0.0	/16
C	192.0.0.0 to 223.255.255.255	192-223 (11000000--11011111 binary)	255.255.255.0	/24
D	224.0.0.0 to 239.255.255.255	224-239 (11100000--11101111 binary)	n/a	n/a
E	240.0.0.0 to 255.255.255.255	240-255 (11110000--11111111 binary)	n/a	n/a

Address Assignment

The following table describes options for assigning IPv4 addresses and other IPv4 configuration values.

Method	Uses
Dynamic Host Configuration Protocol (DHCP)	DHCP is an automatic method for assigning IPv4 address and other TCP/IPv4 configuration parameters to hosts. Client computers contact a DHCP server to receive TCP/IPv4 configuration information. Use DHCP: • For small, medium, or large networks. • For automatic host configuration. • To automatically deliver additional configuration parameters such as default gateway and DNS servers. By default, all Windows computers try to use DHCP for TCP/IPv4 configuration information.

Automatic Private IPv4 Addressing (APIPA)	APIPA is an automatic configuration method where hosts automatically select their own IPv4 address within a specific range. With APIPA: • Windows computers will use APIPA if a DHCP server cannot be contacted. • Hosts select an IPv4 address in the 169.254.0.1 to 169.254.255.255 range with a mask of 255.255.0.0. After choosing the address, the host verifies that no other host on the network is using the selected address. • APIPA sets only the IPv4 address and mask. Because it does not assign a default gateway, APIPA can be used on a single subnet, but cannot be used if communication with other subnets is required. Use APIPA for small, single-subnet networks that do not use DNS servers or do not have Internet or connectivity outside of the local subnet.
Static (manual) assignment	You can manually assign TCP/IPv4 configuration values for a host. • When you configure a static IPv4 address, you must also configure the subnet mask and default gateway. • When you configure a static IPv4 address, you disable DHCP and APIPA. • If you use DHCP you can also assign DNS server addresses manually. Use static addressing: • For small networks that do not often change or grow. • If your network does not have a DHCP server, or if you want to eliminate DHCP traffic from your network. • For specific hosts that must have the same address each time (such as servers). You can use DHCP on the rest of the network and use static addressing for only a few hosts. However, before you use static addressing, explore the possibility of using a DHCP server to assign the same IPv4 address to specific hosts each time an address is requested. • For non-DHCP hosts (hosts that cannot accept an IPv4 address from DHCP).
Alternate IPv4 configuration	With an alternate IPv4 configuration, the system attempts to use DHCP for TCP/IPv4 configuration information. If a DHCP server cannot be contacted, the static configuration values are used. Use an alternate configuration: • If you have a computer (such as a laptop) that connects to two networks: one with a DHCP server and another without a DHCP server.

	• If you want to provide values to properly configure the computer in case the DHCP server is unavailable. When you configure an alternate IPv4 address, APIPA is no longer used.

A Network Address Translation (NAT) router translates multiple private addresses into the single registered IP address.

- The Internet is classified as a *public* network. All devices on the public network must have a unique registered IP address; this address is assigned by the ISP. No two hosts on a public network can have the same IP address.
- The internal network is classified as a *private* network. All devices on the private network use private IP addresses internally, but share the public IP address when accessing the Internet.
- A NAT router associates a port number with each private IP address. Port assignments are made automatically by the NAT router. Communications from the Internet are sent to the public IP address. The NAT router translates the public IP address into the private IP address of the host.
- The Internet Assigned Number Authority (IANA) controls and issues public addressing.
- The private network can use addresses in the following ranges that have been reserved for private use by IANA:

  10.0.0.0 to 10.255.255.255
  172.16.0.0 to 172.31.255.255
  192.168.0.0 to 192.168.255.255

  Internet routers are configured by default to not route private IP addresses.

9.2 IPv4 Custom Addressing

As you study this section, answer the following questions:

- What does a subnet mask identify?
- What is the relationship between CIDR notation and the subnet mask?
- How is a *supernet* different from a *subnet*?
- What are the first and last addresses in a range used for?
- What is the decimal value for a /27 mask?
- How many approximate and actual hosts can you have when using a mask value of /23?
- Given IP addresses and subnet masks, how do you determine if two workstations are on the same subnet?

After finishing this section, you should be able to complete the following tasks:

- Given a network address and custom mask, identify valid subnet addresses.
- Given a scenario with the desired number of hosts, choose a subnet address and mask.
- Given a network address and the subnet mask, identify valid host addresses on that subnet.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  Configure Network Settings for Multiple Subnets
  - Use subnetting to split address ranges
  - Use supernetting to combine address ranges
  - Configure networking for multiple subnets

This section covers the following 70-410 exam objective:

- 401 Configure IPv4 and IPv6 addressing.
  This objective may include but is not limited to:
  Configure IP address options
  Configure subnetting
  Configure supernetting

9.2.1 CIDR and Custom ANDing Basics

CIDR and Custom ANDing Basics
0:00-0:12
Now we're going to talk about CIDRs and custom ANDing basics. The CIDR is a way of expressing a subnet mask and custom ANDing is going to build on our skills we learned in basic ANDing.

Example
0:13-0:39
Let's take a look at some examples. We had talked about the subnet mask in the IPv4 basics and taking a look at some subnet masks in our basic ANDing we had 255s and 0s. So I've got three subnet masks here--one that's just got 255 in the first octet and 0s in the other three octets. This one has two 255s, two 0s, so on and so forth.

CIDR
0:40-2:08
When we take a look at binary we had said, looking at our binary grid, 255 is really just all 1s. The 255 in binary looks just like this; we've got eight 1s. So if we go back to our subnet masks, essentially what they did was say well if I'm looking at a subnet mask like this and I know that 255 is eight 1s; I've got eight 1s here and because I have 32 bits, I have 24 0s. In a subnet mask the 1s are always on the left, the 0s are on the right. But somebody said well instead of having to write out the subnet mask why couldn't we just put a /8 and that would tell everybody that that's eight 1s, one 255, and everything else to the right is a 0 because whatever is not a 1 is going to be a 0. In that case we look at our second subnet mask, that would come out to be a /16; I've got 8 here and 8 here. And then for this one, that would end up being a /24 because I've got 8 here, 8 here, 8 there, so I've got 24 altogether. So this is what they call the CIDR, C-I-D-R. When I started the first book I ever read, it pronounced it cedar; there's lots of people that will pronounce this cider. This stands for the Classless Inter Domain Routing. This is one of those acronyms where you really don't need to know what it stands for if you just know that it's a way of expressing the subnet mask as a / with the number of 1s then you're going to be in great shape.

ANDing
2:09-2:49
So let's take a look at an IP address and some basic ANDing. If this is my IP address, I've got 192.168.3.40/24. That would tell me that I have a 255, which uses up 8; got another 255 that uses up another 8; got my last 255, there's my 24. And so anything that's not a 1 is a 0 and then of course in our basic ANDing we said that anything that's in the same position as a 255 is part of the network ID. So I can do my basic ANDing and come up with a network ID of 192.168.3.0 as the network ID for this IP address. That's pretty easy.

Breaking in the Middle of an Octet
2:50-3:55
Where it starts to get a little bit complicated and where your knowledge of binary is going to come in is when you have something that's not your basic 8, 16, or 24.
So suppose I had a CIDR of /19. Well that's a little bit more difficult. Essentially what's happening is our subnet mask--and I'll go back to the subnet mask for a minute--you know where it stops being 1s and starts being 0s. Everything to the left is my network ID, everything to the right is my host. So here everything to the left is the network ID; everything to the right is the host. When you have something that's not 8, 16, or 24 essentially you're making that break in the middle of an octet. So some of the bits in that octet belong to the network ID, some belong to the host.
Our first challenge is to take that CIDR--we'll go back to the /19--and turn that into a subnet mask. And that's what we're going to do first. And then we're going to see--how can we decide if two IP addresses are local or remote when they're being broken in the middle of an octet. So it's very easy for us to do the basic ANDing. Custom ANDing is a little bit more challenging, but not horrifically so.

Rules of Custom ANDing
3:56-4:32
Now the first thing that we have are some rules. So I'm going to tell you the rules but then we're going to go through them. When we have something that's not 8, 16, or 24 our first step is going to be if the CIDR is greater than or equal to 8, we're going to write down 255 and we're going to subtract 8 from the CIDR. And we're going to do that until we get a number less than 8. At that point we're going to write out as many 1s as the remaining number and we're going to add 0s to get 8 digits. We're going to convert that to decimal and anything remaining to the right is going to be 0. So this is writing down what we're going to do but we're actually going to go ahead and do this.

Example Using the Rules

4:33-5:16

So we'll go back to our /19. So our first step is if the number is greater than or equal to 8, we're going to write down a 255 and then we subtract 8 and that's going to give us 11. Now our second step said we should repeat that until we get less than 8 so again we're still above 8 and I write another 255, subtract 8 and I get 3. As soon as I get a number that's less than 8, no thinking, I just write down that many numbers of 1s. I'll write down that many numbers of 1s, so I'll write down three, one, two, three (111)--don't be afraid to write big--and I'm going to add 0s until I get 8 digits. So if I have three 1s, I've got to add five 0s to get 8 digits because again, these IP addresses are octets; I need 8 digits in binary. So this is what that third octet is going to look like.

Interesting Octet

5:17-6:54

They actually technically call this the Interesting octet, which sounds silly. Sometimes they'll say the subnetted octet but most people say the interesting octet, meaning the one that's not 255 and not 0. So I'm going to take this number and convert it into decimal. And so we would go over to our chart that we use for binary: 1, 2, 3, 1, 2, 3, 4, 5, which means I'm adding up my 128, 64, 32, it's going to give me 14, that's going to give me 12; that should be 224. So my third octet here is going to be 224, right? That's my step three and the last step is anything left over is going to be a 0. There's never going to be more than one interesting octet, so if you come up with more than one interesting octet, something is wrong there. So a /19 gives me 255.255.224.0 and just to make sure, we'll do another one with a /10. 10 is greater than 8; I write down my 255, I subtract 8. At that point I am less than 8, I don't think, I just write down 1, 2, got to add 0s to get eight digits because we need an octet. I add six 0s and if we plug this into our grid I can tell you right now it's going to be 192. Anything left over is going to be 0. A CIDR of /10 would give me a subnet mask of 255.192.0.0.

So far so good so why would we want to break into the middle of an octet? That's a discussion for a different lesson, but in this case we are breaking in the middle of an octet, so that's going to change what happens when we do our ANDing.

Local or Remote?

6:55-7:02

We will go back. We have done basic ANDing, find out if the computers were local or remote just using our regular subnet mask of 8, 16, or 24.

Custom CIDR

7:03-10:10

Now we're going to see what happens when we have a custom CIDR. So here I've got two computers. I've got my 192.168.1.47/27 talking to 192.168.1.80; we want to know if they're local or remote. The first thing I have to do is turn my 27 into a subnet mask. Just as we've done--greater than 8, I'm going to write down my 255; I subtract 8, which is going to give me my 19; still greater than 8, so I've got 255 minus 8 gives me an 11; then another 255 minus 8, so I should get 3 left over; not thinking, 1, 2, 3, fill it out to eight; 1, 2, 3, 4, 5. Some people will stop there and say wait a minute I'm not used to seeing a number in that fourth octet. You want to be very methodical in doing this; don't think. I find when I think, that usually messes me up. I'm just pretending I'm the computer and just following the rules. So if we convert this there--that was our 224--and so now we have our subnet mask and we have our IP address.

Once you've got the subnet mask I recommend you do the easy octets first. So in this case there's always going to be three easy octets. Here we know that anything that's in the same position as a 255 is just itself, so I've got 192. 168 is just going to come right down. The 1 is going to come down. So the only thing we don't know what's going on here is this last octet, the interesting octet. In order to find out what my network ID is, which is my goal, I have to work in binary. I already have the number 224 in binary so what I have to do now is get the number 47 in binary so that I can compare it and do my ANDing out in binary the way the computers do it. So 47, that's going to give me no 128s, no 64s, one 32, give me 15 left over as well; I've got no 16s and 8 with 7 left over. All right, that gives me a 4 with 3 left over, a 2 with 1 left over, and a 1. So this number should be 47 in binary.

Normally I would try to write the 47 above the 224 but it really doesn't matter which one is above or below. What's important is that I've lined up my numbers okay? So I've got my 8 digits roughly together. This is where the term ANDing comes from, where only a combination of a 1 AND a 1 gives me another 1. So what we'll do is just copy these two numbers to a different page so I can have my 47 above my 224 the same way that it appears in the IP address and we get a little bit better perspective on it. So if I look at my first bit here, 0 and a 1, that gives me a 0; a 0 and a 1, that gives me a 0. A 1 and a 1; that's an AND, I get a 1. Two 0s gives me a 0, 1 and a 0 is 0. So since I only have one combination of a 1 and a 1, I'm only going to get one 1 inside the binary number. So this is what my network ID is going to be in the fourth octet.

Converting Back to Decimal

10:11-12:00

What I have to do now is convert that back to decimal and if you convert that back to decimal using your grid, you're going to find out that that should be 32; 128.64.32. So if we go back in to our computer, this particular computer is on the 192.168.1.32 network. Again, it doesn't matter if there's numbers in all four octets if you're not used to seeing that, that's okay.

So now I know that computer A is on 192.168.1.32. Now in order to decide if we're local or remote, I've got to take a look at my other computer; we'll call it computer B. We know that we use the sending computer's subnet mask so I don't have as much work to do here. I already figured out what the subnet mask is. I'll do my easy octets; 192.168.1; so all I really need to know about is what's going to happen here. Very easy to do my 224 because I already know that that's 1, 2, 3, 1, 2, 3, 4, 5; that was my 224 from before and so now all I have to do is figure out 80, which would be no 128s. So I've got a 64, that should give me a 16, okay so no 32s, a 16 and nothing left over. There's 80 and there's 224; again the only thing is to line up the digits and I'm looking for a 1 and a 1 make a 1; everything else is a 0. So in my first one here I'm going to get a 0; there's my 1 and a 1, 0, 0; everything else is going to be 0s and if we convert that back to decimal, that's going to be a 64. So if I'm looking at my computers, this computer's on the 192.168.1.64 network. In that case these two computers are remote because the network IDs are not the same.

Summary

12:01-12:47

We've talked about the CIDR, which is just a way of expressing the subnet mask as a / with the number of 1s, gone through how to take that CIDR and convert it to a subnet mask, and then gone through and done some custom ANDing where we're obtaining the network ID. But now, instead of breaking that portion that's the network and that portion that's the host at a period, we're breaking it in the middle of an octet. We go through, we convert that CIDR to a subnet mask, we do our easy octets, we put the interesting octet in both the IP address and the subnet mask in binary-- only the combination of a 1 AND a 1 is going to give us another 1--and then we find our network ID. Again, if they're the same, we're local; if they're different, we're remote and it's exactly like basic ANDing, just a little bit of binary thrown in to make it more fun.

9.2.2 Subnetting

Subnetting
0:00-0:29
Subnetting comes out of a much earlier time than today. When I first started training, you didn't have private IP addresses-- not that many people had private networks. If you had an IP address, you were on the internet. They only had public IP addresses. And so the issue came up with, if I go out and I get a network, how am I going to make that network that I bought work out for me in terms of usability?

Why Do We Subnet: Binary
0:30-4:14
But before we see how to do subnetting, first, we're going to take a look at some stuff in binary to understand why we would need to subnet, and then we'll go through the process of how to subnet.
Now binary, as we've talked about, is the base 2 number system, and I've just kind of put up a small number here, and you can see that 2^3 is equal to 8. Well, what does it actually look like? I chose something small like 2^3, because I want to see. There's actually 8 different combinations, so by varying the 0s and 1s, I get 8 combinations. These are the actual 8 combinations that I get, so the interesting thing is 2 to whatever number, whatever that power equals, that's the number of combinations that you can get. So, if I had written down 2^4, I would be able to write down a table like this that would give me 16 combinations; 2^5 would give me 32 combinations. So that works out pretty good.
But when I'm trying to look at how many computers I have on a network though, I've got to take out a couple. So if I had a three digit binary number to make IP addresses with, there's two that I can't use. And the two that I can't use are the Network ID, which is where the host is all zeros, and the very last IP address in the network is the Broadcast ID, which is where the host is all ones. So, if I go through and I say, if I had 8 bits and say I had a /24 bit subnet mask, if I'm using 24 bits for my network, then I must have 8 left over for the host. Theoretically, that would give me 256 combinations, but because I can't use my network ID and I can't use my broadcast ID, I'm going to subtract 2, which will give me 254 actual IP addresses that I can use.
So again, when I started teaching, if you had an IP address, you were on the internet. So let's say I have a company, and I went to InterNIC back in you know, 1997 or something like that, and they said, "All right Shad, we'll sell you the network 220.32.17.0." And because that is a Class C network, the default subnet mask for that is a /24, so I would have a /24; that's all the internet routers will recognize. And what that means is all the IP addresses in my company have to start with 220.32.17. I can do whatever I want with the last octet. The only exception being I can't use 0, because that is my network ID--that's the network I bought--and I can't use 255, because that's the broadcast ID. So the IP addresses I can use between there is going to be my 2^8, the 256 minus 2; that gives me 254 usable IP addresses on one network--great, fantastic. The problem is it's one network. Now let's suppose I have a smaller company with three buildings and 50 clients in each building. So I'm well within my 254; I only have 150 clients. The problem is my clients aren't all on one network. They're in three different buildings. So, I haven't got a lot of great choices. Do I go back to InterNIC and buy two more networks and waste all those IP addresses? I've got plenty of IP addresses. No, the better answer is to take this one network that I bought and divide it up into pieces and I'd like to get as close to three pieces as I can.

Creating Pieces
4:15-8:08
Now because of binary I can't get three pieces. What's really going to happen is I'm going to get like four in this case but there's a formula for working it out. But essentially I'm going to take the one network that I bought and I'm going to split it down into pieces. Because I'm splitting it into pieces each piece, each resulting network, is going to be smaller than the one that I bought, but it lets me get more value out of this network that I bought and that's where subnetting came from.
In today's world we don't usually need to subnet that much anymore because there's private networks and there's plenty of IP addresses for private networks; we don't need to do this. Back when I started my career this was a pretty critical skill and is still certainly in use today.
So first we decide how many pieces we've got to split it into. So we said we had three sites. Sometimes they say sites, sometimes they say networks, sometimes they say subnetworks. Those terms all mean the same thing. Your formula is 2 to some power is going to be greater than or equal to number of sites that you need. So we know in our case 2 to some power is going to be greater than or equal to 3. 2^1 gives me two, that's not enough. 2 squared would give me 4 so in this case, my X is going to be 2. But what do I do with that X? That X tells me how many of these 8 bits that belong to me I'm going to borrow to make these new networks that I need. I need 3 new networks. So I'm going to take 2 of those bits that I used to have for hosts and I'm going to turn them into networks. So my 4th octet used to look like this, but I'm taking these first two bits and turning them into networks and I'm going to

get four networks out of that. We saw our 2 squared gives me 4. I'm going to have two 0s, a 0 and a 1, a 1 and a 0, a 1 and a 1, all the different possible--four different possible--combinations I can make out of those 2 bits. The last 6 are all going to be 0s regardless because those are for my hosts. So I'm going to come up with the numbers 0, 64, 128, and 192. Those are the actual network IDs that I'm making by splitting this.
And what does that subnet mask look like? And again I picked a very small number so that we could see it in binary. Normally we don't draw it out like I just did. So once you've got your X equals something, you add that to the default subnet mask. So I've got a default subnet mask of 24, I add my 2 in; my new subnet mask is going to be a /26. That's the CIDR that's going to give me at least three sites or better. So with a /26 we can convert that fairly easy to a subnet mask. So we've got 255, that gives me 18; 255, that gives me 10; 255, that gives me 2. The 2 that I added, 1, 2; 1, 2, 3, 4, 5, 6 would be a 192. So now I've got the network that I bought but I'm going to use a 255.255.255.192 subnet mask which is really a /26, which is going to divide this network up into four pieces; the four numbers that we saw before.
And if we had a lot of networks, like 100, we would never want to go through and draw it out in binary like this. You would be drawing 1s and 0s for the rest of your natural life.

Shortcut to Find Out Network ID's

8:09-9:10
So let me show you a shortcut to figure out what the network IDs are going to be. Now I know I can always have this first network. That's the one I bought. I want to know what the other networks that I've created are going to be. Well the shortcut is to take the interesting octet in the subnet mask, in our case it's 192, and you subtract that from 256. In this case it's going to give us 64. So my networks are going to go in jumps of 64 in that interesting octet, so 0, 64, 128, 192 all right. The next number would be 256, but we can't go above 255, so these are the four networks exactly as I've shown you. We're just using a shortcut to find out what they are.
So now we know our four network IDs; we only need three of them but we've got four. But we're not interested in network IDs and you know we need to find the network ID if we're going to compare if two machines are local or remote.

IP Addresses that can be Given Out

9:11-10:12
We really want to know what IP addresses can we give to these clients such that if somebody were to come along and say hey I'm going to do some ANDing; when you AND them, they're all going to be on the same network because we want to go ahead and set up DHCP or however we're going to assign these we're very much interested in what IP addresses will the actual hosts be given. And we need to have a range such that when we do the ANDing, everybody comes out on their network ID.
Well the first IP address you can give out is just the network ID plus 1. And if I'm doing a grid like this, I do it all the first, all the last, all the broadcast because it's very easy. So here on this network, the first ID I could hand out was .1; 220.32.17; I won't re-write it because my white board's only so big. Here the first one I can give out is 65 and my network ID plus 1. Here the first one I can give out is 129 and here the first one I can give out is 193, that's great.

Last IP Address that can be Given Out

10:13-11:42
Well what's the last IP address I can give out? Well the very last IP address on a network is the broadcast address, which I put a little BC up here for broadcast. So the last IP I can give out is 1 less than the broadcast. So the rule is you take the next network number and subtract 2. So if I'm looking at this very first network here and I want to know what's the last IP address I can give out, my next network is 64, my last IP address that I can give out is 62, my broadcast address is 63 and then here's 64, my next network. The same thing with network number 2 here; the last IP address that I can give out is 128 minus 2 which would be 126, 127 is my broadcast, 128 is the next network. Here my last IP address that I can give out is going to be 192 minus 2, so that gives me 190; 191 is my broadcast, 192 is my next network. Here's the tricky part. The next network would be 256. It's not, because we only have four pieces, but it would have been 256 so we'll subtract our 2 from that. So that very last IP address I can give out here is 254 and then my broadcast is 255.
So that might be a little confusing now but really you just have to go slow, take it step by step and find out what your custom subnet mask is going to be and then your network IDs.

Summary

11:43-12:48
So we went through subnetting, which is the process of taking a network and breaking it down into smaller pieces. We know that we have to stick with the subnet mask that we've been given out on the internet based on the class of the IP address. We find our 2^X is greater than or equal to the number of littler networks or subnets that we need. We add X to that default subnet mask to find our new CIDR--our new subnet mask. We subtract the interesting octet from 256 to find out the jumps that our networks are going in inside of that interesting octet. The first IP address is the network ID plus 1. The last IP address is the next network minus 2. Take your time. Go through it, work through it; as long as you're methodical, you'll get the right answer and you can have just as many networks as you need.

9.2.3 Supernetting

Supernetting
0:00-0:11
This topic is called supernetting, and it's definitely on the Microsoft objectives, but I can't imagine a situation in which you would actually do it nowadays.

Subnetting vs. Supernetting
0:12-0:31
But essentially, if subnetting is the process of taking one network and breaking it down into multiple pieces, and they get more networks but less hosts on each network, supernetting is the process of combining two networks that would have been separate networks into a bigger network.

Example
0:32-0:55
So I've given you a basic class C address, 220.32.17.0, with a default subnet mask, which gives us our standard 254 computers on one network. But here's the scenario: you actually need to get 500 computers. We need to create one network that has 500 hosts. We've got to find some CIDR that's going to accommodate that number of hosts and then find out the associated networks.

Formula
0:56-2:03
The formula is very similar to the formula that we used for subnetting. Subnetting we used 2^X.
Here we're using (2^X)-2, again, because we can't use the network ID and we can't use the broadcast ID. So, (2^X)-2 is going to be greater than or equal to the number of hosts that we need on that network. So, we've got (2^X)-2 greater than or equal to 500. So we want to get as close to 500 as we can; we've got to be over, but we can't be under, but we're trying to get as close as we can. So, I don't have a great way to do this, I just happen to know that 2^7 gives me 128, in that case, 2^8 is going to give me 256, 2^9th is going to give me 512, and there's not even a great way to do this on the calculator-- you've just got to kind of figure it out. So that looks like it's going to be good. So, I could get (2^9)-2 is going to give me 510 computers. 510 is just over the 500 I need, so I'm in good shape, so that tells me that X=9.

CIDR
2:04-2:35
The key is what do I do with that 9? In subnetting, you're adding to the CIDR, so when we're dealing with the number of networks we're adding, and when we're dealing with the number of computers, we're subtracting. And what are we subtracting from? Well, this is telling us the number of hosts that we need. I need 9 bits for the host. I need 9 zeros. Well, if my entire IP address is 32 bits and I need 9 zeros, then that's going to tell me that I need 23 ones. That is my CIDR.

Steps for Supernetting
2:36-2:59
So, our steps for supernetting, (2^X)-2 greater than or equal to the number of hosts that we need, we find out X--in our case, X=9--and we subtract X from 32. That gives us the /23. So, a /23 is going to be sufficient to give us 500 computers or more on one network, and as a subnet mask, it will look like this.

Networks
3:00-3:53
In order to find out what the networks are--we know we have 17.0-- we're going to have to combine some other networks with it. We're going to take the interesting octet in the subnet mask and subtract that from 256, which gives us a 2; in other words, the networks are going in jumps of 2 in that third octet. Well, if my networks are going in jumps of 2 in the third octet, and I've got a 17 in the third octet, the other network that I'm going to be combining is either going to be 16 or 18. And again, I don't know of any math for doing this, you just have to kind of go through and AND, and work it out. In this case, if I looked at that last octet, it looks like there's 2, 3, 4, 5, 6, 7, 8. We're really playing around with this last bit, which could be a 0 or a 1. So that tells me that it should be 16, and 17 should be the same network. But again, I'd have to AND it in order to find out if they're on the same network in real life.

ANDing
3:54-4:39
So I'll show you what that ANDing would look like real quick. We know that 220 comes down, we know the 32 comes down, your 54 is 1, 2, 3, 4, 5, 6, 7. And 17 would be zero 128s, zero 64s, zero 32s, a 16 and a 1. And since we know I can have either a 1 or a 0 in this position--these being the same--that would be the other number, so that would

tell me that I need to be combining 16 and 17. I'm combining them into one big network. Using a /23 is going to put both of these ranges on the same network address.

Summary

4:40-4:48

So that's a brief look at supernetting, which is the process of combining two networks that would have been separate networks into a bigger network.

9.2.4 Custom Addressing Facts

Be aware of the following IPv4 custom addressing concepts.

Concept	Description
Subnetting	*Subnetting* is the process of dividing a large network into smaller networks. When you subnet a network, each network segment (called a *subnet*) has a different network address (also called a *subnet address*).
Supernetting	Supernetting is the process of combining two or more networks. When you create a supernet, you decrease the number of masked bits in the subnet mask. This reduces the number of available subnets, but increases the number of hosts on each subnet.
Classless Addressing	Using custom subnet masks is often called *classless* addressing because the subnet mask cannot be inferred simply from the class of a given IP address. • Using classless addresses is made possible by Classless Inter-Domain Routing (CIDR). • The *CIDR notation* is a syntax for identifying the subnet mask. • The format for the CIDR notation is the slash (/) symbol and the decimal number identifying the number of ones in the subnet mask.
ANDing	ANDing refers to performing a logical AND comparison. You use ANDing to determine the network address of a classless IP address. When determining network addresses, you convert the IP address and the subnet mask to their binary equivalents and compare each bit. The result is 1 when both numbers being compared is 1.

As you work with subnetting operations, use the following tables to quickly find the information you need. By memorizing these tables, you will be able to quickly reproduce the values necessary for identifying the binary and decimal values you use most.

The following table lists the exponent values for powers of 2.

Exponent	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^{10}	2^{16}
Exponent value	2	4	8	16	32	64	128	256	1024	65,536

Memorize the shaded values. To find smaller or larger values, divide or multiply the exponent value by 2. For example, to get the decimal value of 2^{11}, multiply 2^{10} by 2 (giving you 2048). To find the value of 2^{12}, use 2^{10} x 2 x 2 = 4096.

The following table lists the common binary and decimal values used in subnet masks:

Subnet mask value	Decimal equivalent
00000000	0
10000000	128
11000000	192
11100000	224
11110000	240
11111000	248
11111100	252
11111110	254
11111111	255

Use the following table as a shortcut guide to subnetting.

Look for patterns in the table so you can easily reproduce the table at any time.

Masked Bits	Mask Value	Number of Subnets*	Number of Hosts per Subnet	
			Approximate	Actual (2^n - 2)**
/20	255.255.240.0	16	4000	4094

/21	255.255.248.0	32	2000	2046
/22	255.255.252.0	64	1000	1022
/23	255.255.254.0	128	500	510
/24	**255.255.255.0**	**1 or 256**	**256 or 250**	**254**
/25	255.255.255.128	2	128 or 125	126
/26	255.255.255.192	4	64 or 60	62
/27	255.255.255.224	8	32 or 30	30
/28	255.255.255.240	16	16 or 15	14
/29	255.255.255.248	32	8	6
/30	255.255.255.252	64	4	2

*The number of subnets is the number of subnets you get by subnetting a default network address (either class B or class C in this table). For example, if you subnet a class B network using a /24 mask, you would have 256 subnets.
**To identify the actual number of hosts per subnet, use the formula $2^n - 2$, where n is the number of unmasked bits in the subnet mask. Remember to subtract 2 for the addresses that are not assigned to hosts:

- The first address in the range is the subnet address and cannot be assigned to hosts.
- The last address in the range is the broadcast address and cannot be assigned to hosts.

To discover if workstations are on the same subnetwork, perform the following calculation:

1. Calculate the binary value of the subnet mask and determine which octet is affected by the subnet mask. For example, a /26 subnet mask affects the last octet as shown below:

   XXXXXXXX.XXXXXXXX.XXXXXXXX.XX000000 = /26

2. For the affected octet, determine how many subnets are available within the subnet mask and calculate the decimal value for each subnet. For example, a /26 subnet mask has four subnets available as shown below:

.00000000 = .0
.01000000 = .64
.10000000 = .128
.11000000 = .192

3. For the affected octet, remove the first IP address (network address) and last IP address (broadcast address) in the subnet(s) as possible host addresses. For example, a /26 subnet mask the possible IP addresses for the first subnet are:

   .00**000000** = .0 (Network address for first possible subnet. This address is *not* valid for a workstation IP address.)
   .00**000001** = .1
   .00 **000010** = .2
   ...
   .00**111110** = .62
   .00**111111** = .63 (Broadcast address. This address is *not* valid for a workstation IP address.)
   . **01**000000 = .64 (Network address for the next possible subnet.)

4. Determine if all of the assigned IP addresses fall within the same subnet. For example, the first possible subnet for a /26 subnet mask could have IP addresses in the .1 - .62 range:

   .00**000001** = .1 (valid address)
   .00 **000010** = .2 (valid address)
   ... (valid addresses)
   .00**111101** = .61 (valid address)
   .00 **111110** = .62 (valid address)

9.3 IPv6 Addressing

As you study this section, answer the following questions:

- How does IPv6 support differ on various Microsoft operating systems?
- What limitations are associated with ISATAP in an IPv6 implementation?
- Which IPv6 tunneling methods work through NAT?
- When should you implement Teredo?
- When is 6-to-4 tunneling automatically configured?
- What technology allows an IPv4-only host to communicate with an IPv6-only host?

After finishing this section, you should be able to complete the following task:

- Configure IPv6 settings.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  - Configure Basic Network Settings
    - Configure IPv6 Settings
  - Configure Network Settings for Multiple Subnets
    - Configure Networking for Multiple Subnets

This section covers the following 70-410 exam objective:

- 401 Configure IPv4 and IPv6 addressing.
  This objective may include but is not limited to:
  - Configure interoperability between IPv4 and IPv6
  - Configure ISATAP
  - Configure Teredo

9.3.1 IPv6 Basics

IPv6 Basics
0:00-1:30
We're going to go over the basics of IPv6. IPv6 was created because, even very early in the history of the Internet, IPv4 was running out of addresses. They wanted to create a new system of IP addresses that would go ahead and allow for expansion for the internet to become really as big as they could possibly imagine it would be. Let's take a look at some of the rules and take a look at how the addresses work. IPv6 follows the same rules as IPv4. Every host on the network must have a unique IP address. All the hosts on the same network must have the same network ID. We're going to use a little bit different language in a minute. Hosts can only communicate directly with other hosts on their own network. If I'm going to communicate with a host that's not on my network, I've got to go through a router. The difference with IPv6 addresses is, instead of being 32 bits like IPv4, they're 128 bits. With 128 bits, it's a lot of addresses. I haven't memorized exactly how many, but I once read it was something like 667 sextillion addresses for every square meter of the earth's surface. That's a lot, though we shouldn't be running out of these addresses anytime soon. They are expressed in eight blocks of 4-digit hexadecimal numbers separated by colons.

Example IPv6 Address
1:31-7:30
Here's an example of an IPv6 address. You can see that I've got eight blocks. This is actually much easier to write. If we had to write the whole 128 bits, it'd be a lot longer. Let me just briefly go in and show you hexadecimal. You do not need to know how to convert back and forth between hexadecimal. Sit back, relax, just enjoy the tour. Let's take a look at an example of hexadecimal. Hexadecimal is base 16. In the first column we have 16 to the zero, which is just 1. 16 to the 1, which is just 16. 16 squared, which is 256. I went all the way up to 16 cubed which is 4096. The problem comes in because, as we know with number systems, we can have just one digit in each position. Zero through nine is no problem, but if I get to 10, I can't have two digits in that column. Ten can be represented by the letter A, 11 is B, 12 is C, 13 is D, 14 is E, and 15 is F. I threw together a simple hexadecimal number, A2E8. A or 10 x 4096 gives me 40960, 2 x 256 gives me 512. E is 14, so 14 x 16 gave me 224, and then 8 x 1 was just 8. We add them all together, and we get 41704. You can see these numbers get really big really fast. You don't need to know hexadecimal. Hexadecimal is used for IPv6 and MAC addresses. The moral of the story is this: you're never going to see any letters other than A through F. If somebody asks you to take a look and say what's a good IPv6 or MAC address, don't pick G, don't pick K, don't pick Z. It's just going to be the letters A through F that you're going to see in these types of addresses. When I saw my first IPv6 number, I was horrified. I thought, "Oh my gosh, subnetting with IPv4 is very difficult. How am I going to handle this?" Remember, they were inventing IPv6 after the internet already existed. Now, they were looking at it as a global network. What they really wanted to do was make it easier to set up networks using IPv6. If an IPv6 address is 128 bits, normally the network portion is 64--the first half. Now we know from our rules of TCP/IP that hosts can only communicate directly with hosts on their own network. All the computers on the same network must have the same network ID, or an IPv6, the same prefix. That includes the network card of the router that's on that network. It gets really slick; you can program your router to advertise the prefix or the network ID of that network. Your IPv6 client boots up. It listens. It actually doesn't use broadcast. It uses multicast, but it will receive that announcement and say, "Hey, I'm on network blah, blah, blah, 64 binary numbers." Now it can automatically make up the second half of that address. 64 bits are being used for the network ID. The host part of that is going to be another 64 bits. Subnetting can be done using the same formulas as IPv4. There's no need to subnet IPv6. We can really stick with this 64 bit network ID and be very happy for the rest of your life. God forbid somebody asks you, just use the same exact formula. The routers can be configured to announce the prefix, so the computer will know exactly what network it has booted up on. The clients can use either a randomly generated host address. They just make up a 64 bit number and then make sure nobody else is using it. Or--this gets even more slick-- they can use the EUI-64. Every network card that's manufactured has a unique address burned into the network card. This is called the MAC address. It's a 48 bit hexadecimal number for IPv4. The first half of it was three sets identifying the manufacturer, the last three identify that particular network card. What's incredible is the manufacturers are organized. They have chunks of numbers where they literally can guarantee that every network card in the world has a unique MAC address. Network cards that work with IPv6 have an IPv6 equivalent MAC address that's 64 bits. That's your EUI-64. We can actually program it so that the computer will get that prefix from the router, tack on this IPv6 equivalent MAC address, the EUI64. Bingo. I've got a good address. I know it's unique, because I know that EUI-64 is unique, because the manufacturers are ensuring that. Now I've got a unique address for the host on that network. I know that it has the same network ID as everybody else on that network. I know it can communicate with a router, because the router has the same network ID, because that's where I got my network ID.

It virtually eliminates the need for DHCP. The only thing we would need DHCP for with IPv6 is to hand out additional information, like who's the DNS server on this network, or there's WDS, or some other service that's in play that we can't get automatically. It's really streamline. Once you get used to it, it's actually pretty cool.

Global-Unicast Addresses

7:31-8:06

The last thing we need to talk about are IPv6 address types. To make it easier, I've listed the IPv4 address types that you might be familiar with. Public addresses are addresses out on the internet. We don't say public addresses in IPv6. We say global unicast. If you see global unicast, that just means an IPv6 address out on the internet. They start with a two or three. If the very first digit is a two like the example I showed you, you immediately know that that's an address out on the internet. It's public address. In IPv4, we also have private addresses.

Unique-Local Addresses

8:07-8:34

These are addresses that have been pulled off the internet for use in private companies. You know you need a private address because you have internal routers. In IPv6, these are called unique-local addresses, and they start with an FC or an FD. Originally, when they first came out, they had a little bit different terminology for the private addresses. It wasn't unique-local. I believe it was site-local.

Link-Local Addresses

8:35-9:04

In IPv4, we have APIPA addresses, automatic private IP addresses. This is used whenever DHCP is broken. If I can't talk to DHCP but I'm trying to use to a DHCP address, I used an APIPA address. In IPv6, we called this link local addresses, and they start with FE8. You do an ipconfig /all on any computer from Vista Windows Server 2008 on up, you're going to see an FE8 address, because IPv6 is running by default.

Multicast Addresses

9:05-9:53

Last type is multicast addresses. They are also called multicast addresses. In IPv6, they start with an FF. Unicast means one computer talking to one other computer. The complete opposite of that would be your broadcast. It's one computer talking to everybody. We were standing in a classroom, I consider what I do teaching as broadcasting, which means I'm talking to everybody, and whoever thinks they're being addressed can respond. Multicast splits the difference. The packet is going to be addressed to a number of hosts, but not everybody on the network. I send that packet to a multicast address, and behind that is almost sort of a distribution list of IP addresses of the host that will receive that packet.

Summary

9:54-10:31

That's pretty much the basics of IPv6. There are 128 bit addresses. They are expressed in eight blocks of hexadecimal, hexadecimal numbers only having the letters A through F. We can drop leading zeros. We can express blocks of zeros with a double colon. We can only use the double colon once. We have different types and classes. We know exactly what type of IPv6 address it is just by looking at the first few digits. You want to make sure that you have those memorized. That should get you pretty far with IPv6 so that you can get started and get used to this new technology, which is pretty cool.

9.3.2 Configuring IPv6

Configuring IPv6
0:00-0:33
In this video, we're going to take a look at configuring IPv6. There's a couple of ways to get into the right dialog box. One way would be to go to Local Server, and I can just click on this Ethernet here, and that will show me my network adaptor. If for some reason you don't like that, another easy way to get in is to come down here into the bottom right hand corner, right click the icon for the Network, and Open Network and Sharing Center. I would then Change Adaptor Settings, and once I'm at my adaptor, I'm going to right click it and go to Properties.

IPv6 Properties
0:34-0:40
I'll go into the Properties of IPv6. By default, it's set up to Obtain an IP address automatically.

Typing in an IPV6 Address
0:41-0:54
If I wanted to, I could type in an IPv6 address.
IPv6 addresses are made up of eight blocks of four digit hexadecimal numbers. By default, it chooses a subnet prefix length of 64 because, by default, half of the IP address is the network ID, the other half is this particular host.

Subnet Prefix Length
0:55-1:10
I could put in a default gateway if I need to, and I can put in an IPv6 DNS server.

Advanced TCP/IP Settings
1:11-1:27
In Advanced, I can add multiple IPv6 addresses, and if I need to, I can add multiple default gateways. That would be used if there were a backup router in my environment. Make sure if you add a second default gateway as a backup you set the metric of the backup to be less than the primary.

Advanced DNS Configuration
1:28-1:43
I can also do some advanced DNS configuration. I can add multiple DNS server addresses. I can set up which DNS suffixes would be appended to single label names, and tell the computer whether or not to register this connection's address in DNS.

IPv6 and IPv4 Details
1:44-2:12
If you want to check your IPv6 information, or IPv4, for that matter, there's a couple of great ways to do that. I can right click and go to Status and then hit Details. That will show me my IPv6 information and my IPv4 information.
If I didn't see IPv6 information or IPv4 information in this dialog box, it would be because those checkboxes aren't checked in the properties of the adaptor. That's how we configure IPv6.

9.3.3 IPv6 Facts

Because of the rampant Internet growth, IPv4 addresses are being depleted. Many organizations use Network Address Translators (NATs) to map multiple private address spaces to a single public IP address. However, using NATs to overcome the problem introduces security related issues as well as other problems when connecting two organizations that use the same private address space. The IPv6 address standard seeks to address the issues of the IPv4 address standard.

IPv6 follows the same basic rules as IPv4:

- Each host must have a unique IPv6 address.
- Each host on the same logical network must have the same network address.
- Hosts can communicate directly only with other hosts on the same logical network.

The following table describes the structure of an IPv6 address:

Component	Description
Format	IPv6 uses a 128-bit address made up of 32 hexadecimal numbers, organized into 8 quartets. • The quartets are separated by colons. • Each quartet is represented as a hexadecimal number between 0 and FFFF. Each quartet represents 16-bits of data (FFFF = 1111 1111 1111 1111).
Leading zeros	Leading zeros can be omitted in each section. For example, the quartet 0284 could also be represented by 284. • Addresses with consecutive zeros can be expressed more concisely by substituting a double-colon for the group of zeros. For example: FEC0:0:0:0:78CD:1283:F398:23AB FEC0::78CD:1283:F398:23AB (concise form) • If an address has more than one consecutive location where one or more quartets are all zeros, only one location can be abbreviated. For example, FEC2:0:0:0:78CA:0:0:23AB could be abbreviated as either one of the following: FEC2::78CA:0:0:23AB FEC2:0:0:0:78CA::23AB FEC2::78CA::23AB is not a valid IPv6 address.
Prefix and interface ID	The 128-bit address contains two parts: • The *prefix* is equivalent to the IPv4 network ID. The prefix is typically 64 bits.

	IPv6 addresses are allocated based on physical location, with the prefix also including the global routing information. The 64-bit prefix is often referred to as the *global routing* prefix. Routers can be configured to announce the prefix. Subnetting can be done using the same formulas as IPv4. • The last 64-bits make up the *interface ID*. This is the unique address assigned to an interface. Clients can use a randomly generated host address or the Extended Unique Identifier 64 (EUI-64). Addresses are assigned to interfaces (network connections), not to the host. Technically, the interface ID is *not* a host address.

When using IPv6, be aware that:

- The IPv6 loopback address is ::1
- IPv6 eliminates the need for DHCP to issue IP addresses. DHCP is needed only to provide additional information, such as the DNS server information.
- IPv6 makes it possible for each device to have a publicly registered address. Having a unique address for each device removes the need for NAT and PAT.

The following table compares IPv6 address types with IPv4 address types:

IPv4 Address	IPv6 Address	IPv6 Starts With
Public Addresses	Global-Unicast	2 or 3
Private IP Addresses	Unique-Local	FC or FD
APIPA	Link-Local	FE8
Multicast	Multicast	FF

Keep in mind the following about addressing:

- A Global-Unicast address is an address on the Internet.
- Unique-Local, previously referred to as Site-Local, indicates a private IP Address.
- Link-Local indicates that the IP address was configured by default.
- Multicast indicates that the packet is addressed to a number of hosts on the network, but not all hosts.

ISATAP and Teredo
0:00-0:03
In this video, we're going to talk about ISATAP and Teredo.
Background
0:04-0:26
Basically, the background is this: if you want to use IPv6, you have to support IPv6 all the way from the sender to the receiver. As IPv6 has been developed, it's difficult to set that infrastructure up. It's almost always the case that, at some point, these packets are going to have to cross an IPv4 network, unless your company has gone to native IPv6.
Tunneling
0:27-0:40
There are a number of technologies for having an IPv6 packet cross an IPv4 network, and generally it's referred to as tunneling. It's a whole bunch of them. We're going to talk about ISATAP and Teredo.
First of all, ISATAP basically encapsulates your IPv6 traffic in an IPv4 packet.
ISATAP
0:41-1:01
When I think of this, I think of this as sending internal mail in a big company. I have my letter, it's addressed just fine, but then I stick it in one of those inter-office envelopes and that's how it travels to the other office, and then it gets taken out and put into whoever's box.
Designed for Intranet Usage
1:02-1:16
The key with ISATAP is it's designed for intranet usage, so it's not really supposed to be for the Internet. It definitely does not work through NAT, so that's a limitation. It does require an ISATAP compatible router, which you'll have to set up.
ISATAP Addresses
1:17-1:51
ISATAP uses the link local addresses, so it uses the FE80 with a 64 prefix. For the host portion of the address, we have a block here of four zeros, and then you're going to see 5EFE, and then the last two blocks will be your IPv4 address in hexadecimal. A good trick is to know ISATAP addresses always start with FE80:: a double colon, because I can get rid of any zeros, and then 5EFE:, and then some more numbers for the address. If it starts like this, it's an ISATAP address. I can configure ISATAP using netsh.
Teredo
1:52-1:58
Teredo is another technology, and they all do the same thing. They encapsulate IPv6 in an IPv4 packet.
Designed for Internet Usage
1:59-2:10
They key with Teredo is it is designed for Internet usage, and it can work through NAT. Really, all I have tucked in my mind about Teredo is IPv6, through IPv4, behind NAT. You know that, I think you'll be pretty good.
Teredo Addresses
2:11-2:33
It uses Global-Unicast addresses, which means they all start with 2001 and a 32-bit prefix. What happens is it actually adds the IPv4 address in hexadecimal to come up with the total prefix, so I'm going to see 2001, zero, then probably a couple of blocks; that's my IPv4 address in hexadecimal, slash 64 for Teredo.
Direct Access
2:34-3:51
Why do we care about all this? A lot of it happens in the background, and unless you're really transitioning to IPv6, probably not going to be involved. The reason it's really important is because some of these technologies are required to support Direct Access. Direct Access is a VPN replacement. Microsoft would say, "Let's talk about the bad things about VPN." The user has to double-click a VPN icon; they're probably only going to do that when they want to get their e-mail. Group Policy doesn't come down. If they get disconnected, they have to get reconnected, they can't make the icon, blah, blah, blah.
The idea behind Direct Access was as soon as the client boots, it contacts a Web server called the Network Location Server, to find out if it's inside the company or outside. If it can't contact the NLS server, it must be outside the company, and it immediately initiates a direct access connection, assuming it has internet access. What that means is your work computers--assuming they have internet access--are going to be on the work network before the user

even logs in. If they log in, and then manually connect to a network, like they're at a hotel, if you connect to wireless, the minute they get internet connectivity, direct access is going to connect them to the work network, which means Group Policy will come down, I can look at what they're doing, I can enforce any rules, I can really control that computer.

Direct Access and Teredo

3:52-4:07

Direct Access used to be very heavily dependent on ISATAP. They're moving away from that. It is possible that your Direct Access setup might require Teredo. We're just going to take a look at a couple of facts about Teredo with Direct Access, and then we'll be all set.

Test Teredo

4:08-4:30

If you need to test Teredo, you're a Direct Access client, you could actually just do an ipconfig /all. And what you're looking for is a section that says, "Tunnel adapter Teredo Tunneling Pseudo-Interface". That's actually what it's called, and then it should have a description, "Microsoft Teredo Tunneling Adapter", and you should see an IPv6 address that starts with 2001, because that's what Teredo addresses start with.

Enable Teredo

4:31-5:24

If you don't see that, you need to enable Teredo, and here's the command to do that: netsh interface Teredo, I'm going to set the state of Teredo to an enterpriseclient. If you want to test and make sure it's working okay, you can use this PowerShell command to find out who your direct access DNS server is, and if you can ping that DNS server, using its IPv6, then Teredo is working.

In terms of ISATAP, Microsoft recommends using NAT64 instead of ISATAP. Direct Access was really a lot of work to set up in Windows 2008 R2, and one of the things that caused problems was ISATAP, because the DNS servers, by default, don't allow registrations from ISATAP clients, so you used to have to run special commands to open it up. Now they've really streamlined Direct Access, so that you're able to run a wizard and everything's going to be configured the way it needs to be configured. That's how you test Teredo, if you need to do that.

9.3.6 IPv4 and IPv6 Interoperability Facts

Transitioning to IPv6 requires time and dedication. IPv6 is not backwards compatible with IPv4: IPv4 hosts and routers do not support IPv6 traffic, and IPv6 hosts and routers do not support IPv4 traffic.

The following table lists various strategies for deploying IPv6:

Method	Description	
Dual stack	A common method for moving from IPv4 to IPv6 is referred to as *dual stack* configuration. In this method, both the IPv4 and IPv6 protocol stacks run concurrently on a host. IPv4 is used to communicate with IPv4 hosts, and IPv6 is used to communicate with IPv6 hosts. Microsoft uses two methods to create a dual stack host: • Windows 2003/XP uses a dual stack implementation, where IPv4 and IPv6 are separate protocols. • Windows Vista and later, as well as Windows Server 2008 and later, use a dual architecture protocol stack, where IPv4 and IPv6 use common transport and framing layers. By default, Windows uses IPv6 whenever possible. The dual layer architecture means you cannot uninstall either IPv4 or IPv6; however, you can disable one or the other, or change their order of priority.	
Tunneling	*Tunneling* wraps an IPv6 packet within an IPv4 packet, allowing IPv6 hosts or sites to communicate over the existing IPv4 infrastructure. With tunneling, a device encapsulates IPv6 packets in IPv4 packets for transmission across an IPv4 network, and then the packets are de-encapsulated to their original IPv6 packets by another device at the other end. You can configure the following tunnel types, and tunnels can be configured manually or automatically: • Router-to-router • Host-to-router or router-to-host • Host-to-host (end-to-end) Windows Server 2008 and later and Windows clients support the tunneling solutions listed below.	
	Manually configured tunnel	With a manually configured tunnel, tunnel endpoints are configured as point-to-point connections between devices. Manual tunneling: • Is configured between routers at different sites.

		• Requires dual layer routers as the tunnel endpoints. Hosts can be IPv6-only hosts. • Works through NAT. • Uses a static (manual) association of an IPv6 address with the IPv4 address of the destination tunnel endpoint. • Is configured using **Netsh**. Because of the time and effort required for configuration, use manually configured tunnels only when you have a few sites that need to connect through the IPv4 Internet, or when you want to configure secure site-to-site associations.
	Intra-site Automatic Tunnel Addressing Protocol (ISATAP)	The Intra-site Automatic Tunnel Addressing Protocol (ISATAP) is a tunneling method for use *within* a site to provide IPv6 communication over a private IPv4 network. ISATAP tunneling: • Is configured between individual hosts and an ISATAP router. • Requires an IPv6 router to perform tunneling, and dual layer or IPv6-only clients. Routers and hosts perform tunneling when communicating on the IPv4 network. • Does *not* work through NAT. • Automatically generates link-local addresses that includes the IPv4 address of each host: The prefix is the well-known link-local prefix: FE80::/16. The remaining prefix values are set to 0. The first two quartets of the interface ID are set to 0000:5EFE. The remaining two quartets use the IPv4 address, written in either dotted-decimal or hexadecimal notation. A host with an IPv4 address of 192.168.12.155 would have the following IPv6 address when using ISATAP: FE80::5EFE:C0A8:0C9B (also designated as FE80::5EFE:192.168.12.155). Use ISATAP to begin a transition to IPv6 *within* a site. • You can start by adding a single ISATAP router and configuring each host as an ISATAP client.

		• Vista clients will use ISATAP automatically if they can find the ISATAP router. • Vista clients query the DNS server for a router named ISATAP. When using ISATAP, be sure to use this name for the server, or create an A or CNAME record in DNS using ISATAP as the name and pointing to the ISATAP router.
	6-to-4 tunneling	With 6-to-4 tunneling, tunneling endpoints are configured automatically between devices. 6-to-4 tunneling: • Is configured between routers at different sites. • Requires routers that provide dual layer support as the tunnel endpoints. Hosts can be IPv6-only hosts. • Works through NAT. • Uses a dynamic association of an IPv6 site prefix to the IPv4 address of the destination tunnel endpoint. • Automatically generates an IPv6 address for the site using the 2002::/16 prefix followed by the public IPv4 address of the tunnel endpoint router. For example, a router with the IPv4 address of 207.142.131.202 would serve the site with the following prefix: 2002:CF8E:83CA::/48 (CF8E:83CA is the hexadecimal equivalent of 207.142.131.202). Use 6-to-4 tunneling to dynamically connect multiple sites through the IPv4 Internet. Because of its dynamic configuration, 6-to-4 tunneling is easier to administer than manual tunneling.
	Teredo tunneling	Teredo (also known as NAT traversal or NAT-T) establishes the tunnel between individual IPv6 hosts so they can communicate through a private or public IPv4 network. Teredo is a last resort technology in that it is used only when there is no native IPv6, ISATAP, or 6-to-4 connectivity present between hosts. Teredo tunneling: • Is configured between individual hosts. • Has dual layer hosts that perform tunneling of IPv6 to send on the IPv4 network. • Works through NAT. • Uses a 2001::/32 prefix followed by the IPv4 public address converted to hexadecimal. For example, the IPv4 public address of

		207.142.131.202 would provide clients with a prefix of 2001:0:CF8E:83CA::/64. For Windows Vista and Windows 7, the Teredo component is enabled but inactive by default. In Windows Server 2012, Teredo is enabled by default only on non-domain networks (it is disabled by default on Windows Server 2008 and 2003 SP1). To use Teredo, a user must either install an application that needs to use Teredo, or configure the advanced settings on a Windows Firewall exception to use edge traversal. Teredo behavior differs when machines are members of a domain. Teredo is disabled on XP and Server 2003 machines that belong to a domain. Teredo is enabled on Vista and 2008 machines that belong to a domain. Teredo is disabled by default on Windows 8 and Windows Server 2012 machines that are part of a domain.
PortProxy	PortProxy is a TCP proxy that allows an IPv4-only host to communicate with an IPv6-only host. PortProxy does this by transmitting TCP traffic for application-layer protocols that do not embed address or port information in the TCP segment. Thus, an application like FTP does not work across a PortProxy computer because FTP embeds addresses when using the FTP Port command. To configure PortProxy, use the **Netsh interface portproxy** command with the necessary parameters.	
IPv4 Compatible Address	An IPv4 address that is compatible with IPv6 has ten octets, with the last four octets as the IPv4 address of the device. The format is: 0:0:0:0:0:0:w:x:y:z	
IPv4 Mapped Address	If a device is not compatible with IPv6, you can use an IPv4 mapped address. This address is used to represent an IPv4 only node to an IPv6 node. The sixth octet contains FFFF with the last four octets as the IPv4 address of the device. The format is: 0:0:0:0:0:FFFF:w:x:y:z ::FFFF:w.x.y.z is a simplified version.	
IPv6 to IPv4 Address	An IPv6 to IPv4 address allows IPv6 packets to travel over an IPv4 network, such as the IPv4 Internet, without additional configuration or tunneling. This type of addressing works best when an IPv6 to IPv4 router is used. The first octet is 2002, the second octet contains the first two bytes of the IPv4 address, and the third octet contains the second two bytes of the IPv4 address. The format is: 2002:u:v::/16	

10.1 DHCP Basics

As you study this section, answer the following questions:

- What are the steps a client uses to acquire an address from a DHCP server?
- When must you authorize a DHCP server? What permissions do you need to authorize a DHCP server?
- Why does a DHCP server shut down if its address is not found in Active Directory? What does this protect against?
- How can you change the subnet on a scope?

After finishing this section, you should be able to complete the following tasks:

- Install and authorize a DHCP server.
- Create and activate scopes.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  - Installation and Authorization
    - Install the DHCP Role
    - Authorize DHCP Servers
  - Manage DHCP Scopes, Exclusions, and Reservations
    - Create and Configure Scopes

This section covers the following 70-410 exam objective:

- 402. Deploy and configure Dynamic Host Configuration Protocol (DHCP) service.
  This objective may include but is not limited to:
  - Create and configure scopes
  - Configure DHCP options
  - Authorize DHCP server

10.1.1 DHCP Overview

DHCP Overview
0:00-0:51
In this video, we're going to go through an overview of DHCP. DHCP stands for Dynamic Host Configuration Protocol, but I usually don't try to memorize that. What I know about it is it hands out IP addresses. The idea behind DHCP is, we don't want to visit each and every client. Our focus is centralized administration. You give me 10 computers and ask me to type in an IP address in all 10 computers, I will guarantee you that at least one or two of them aren't going to work. With DHCP, the computers can obtain an IP address from the DHCP server, which allows me to centrally manage my IP addresses. It also allows my clients to be very portable, because if they move from network to network, they can get an IP address that's appropriate for whatever network they're connected to.
We're going to go through the process that the client uses to obtain an IP address and talk about some of the politics of just getting the DHCP up and running on a very basic level.

The Client Obtains an IP Address
0:52-1:00
Here I have a very basic situation. I've got my client, it's connected to a switch, and there's a DHCP server on the same switch. We'll go ahead and give them network cables.

DHCP Discover Broadcast
1:01-1:41
When the client boots up, it sends out a broadcast called a DHCP Discover Broadcast, basically saying, "help, is there a DHCP server in the house?" Every DHCP server that gets that broadcast will send out a broadcast in response, pretty much giving an offer that essentially says, pick me, pick me, pick me. All the DHCP servers will respond. The client is going to request an IP address from the very first DHCP server that it gets an answer from. It sends out another broadcast saying, "okay, DHCP, I pick you". The DHCP server sends out a last broadcast acknowledging that it's gotten the request and giving the client the IP address.

DORA Process
1:42-2:11
This process that the client goes through to obtain the IP address is called the DORA process. You don't have to know that to function with DHCP. The most important thing about this process is that it's done with broadcasts. It has to be done with broadcasts because the client doesn't have an IP address, and it certainly doesn't know the IP address of the DHCP server. It comes up, "help, is there a DHCP server in the house?" They all respond. DHCP1, I pick you. Great, you've got your IP address, and then after that, we don't need any more broadcasts.

DHCP Discover Broadcast Across a Router
2:12-3:00
What's critical about the broadcast is this: if you have a situation where the DHCP server is across a router, then you're going to have some problems. Here, I have my client connected to a switch, here's a router, and it's going to start my DORA process by sending off the DHCP Discover packet, which goes out as a broadcast. Routers do not pass broadcast traffic. Routers, by definition, are connected to two or more different networks. Their job is to pass information intended for the other network. Well a broadcast is intended for everybody on this network, but not anybody else.
The router looks at that and says, "if the client is on Network A, the DHCP server is on Network B." A broadcast on Network A, by definition, is not intended for Network B, then the router won't do anything with it. If you do have a situation like that, you're going to have to jump through a few extra hoops in order to make DHCP work.

Rogue DHCP Server
3:01-4:11
The other important thing to know about the DORA process is that the computer picks the first DHCP server that it gets a response from. In an ideal situation, it's probably not going to get a response from more than one DHCP server. If somebody creates a DHCP server in your environment that you're not aware of, we call that a rogue DHCP server. If I really wanted to mess up your network, I could bring in a little home router from home that's also a DHCP server, plug that in. Some of the clients are going to get IP addresses from your DHCP server, some of them are going to get IP addresses from my DHCP server, and you're going to be running around troubleshooting network IP address problems all day long until you catch me and I get fired. To protect from Microsoft servers being used for this purpose, maybe not even maliciously but inadvertently, if Active Directory is in the environment, they can detect Active Directory, and if they're not a member of the domain and they haven't been authorized as DHCP servers in that domain, they'll simply shut down. They won't give out IP addresses. Even if you have a member server that's a DHCP server, it checks for the presence of Active Directory about every five minutes. As soon as it detects Active Directory, if it's not authorized, that's it. It will stop giving out addresses.

Active Directory

4:12-4:46

Just keep in mind, if you have Active Directory in the environment, your DHCP server must be a domain member. If it's not working, you have Active Directory and it's not a member of the domain, that might be your problem. It also needs to be authorized within Active Directory to be allowed to give out IP addresses, so that could be your problem as well. Maybe it's not authorized. In order to authorize it, you must be a member of Domain Admins. You can also be a member of Enterprise Admins, but Domain Admins at least. If you're logged on as a member of Domain Admins, when you install DHCP, sometimes the DHCP server gets automatically authorized. That's something I always check right off the bat.

Summary

4:47-5:04

DHCP is a centralized solution for distributing IP addresses to clients. It allows me to control the IP addresses so that clients get the correct IP address when they connect to the network. It's based on a series of broadcasts, and the DHCP server must be a member of Active Directory and authorized to make sure that it's not a rogue DHCP server.

10.1.2 DHCP Scopes and Options

DHCP Scopes and Options
0:00-0:06
Let's talk about the DHCP scope and options. Once you have your DHCP server installed and authorized, the next step is to create a scope.

Scope
0:07-0:35
The scope is the range of IP addresses that can be handed out by the DHCP server. That scope has a particular subnet mask associated with it, which means it's a scope for that network, that network ID. There can only be one scope per network.
If I have a scope for the 192.168.1.0 network with a /24 subnet mask, I can only have one scope for that network. If I need to limit the IP addresses given out, that can be done a little bit differently. There's just one scope per subnet.

Option
0:36-1:03
If I need to hand out any extra information with the IP address, that's done as an option. Here we have the scope--remember that's the range of IP addresses. There's one scope per subnet. The scope itself actually needs to be activated before it will work, but that will all be prompted by the wizard.
Options are extra information that I'm going to give out with the IP address and subnet mask. There are four levels of options that we can give out.

Overview of Option Types
1:04-1:40
Server options affect everyone who is a client of that server, regardless of which scope they obtain an IP address from. Scope options are given out to everyone who's a client of that particular scope. Class options are given to everyone who is a member of that class, and you need to set the class at the actual machine so it knows what class it is. The last ones are client options. Client options are given out to just that particular client. It's possible that if I'm using a centralized DHCP solution, that my DHCP server may have multiple scopes. If I have two networks, A and B, I would have a separate scope for each of those networks.

Server Options
1:41-1:59
Server options are options that are given to anyone who obtains an IP address from that server. A great example of this might be DNS. Maybe I only have one DNS server for my entire network, it's located on network B, but it doesn't matter if the client's on network A or network B, they're all going to use the same DNS server. I can implement that as a server option.

Scope Options
2:00-2:23
Scope options only apply to clients of that scope. Default gateway, or the address of the router, is going to be different depending on which network I am on, because my router has to be on the same network as the client. That's something that would always be done as a scope option. If I get an IP address from scope A, I'm going to get one type of default gateway, if I get it from scope B, I'm going to get a different default gateway. Whatever is appropriate for my network.

Class Options
2:24-2:40
Class options are very rarely used. If you do have a situation where only HP computers or some class of computers need a particular option, you would run a command on the client to set the class, and then you could create a class option in DHCP. Only clients of that class will receive that option.

Client Options
2:41-2:50
The last ones are client options, and those are specific to a particular client. They're actually done in a reservation, which we'll talk about in a different video. They apply only to that particular client.

MAC Address
2:51-3:08
How does the DHCP server know that that's the actual client who should get that option? Every network card in existence has a unique identifier called the MAC address. When we set up our client options, we're going to use the MAC address to identify that particular client as being the only one that's going to get that option.

Summary

3:09-3:27

With DHCP, we create a scope, which is the range of IP addresses that the DHCP server can give out to clients. Along with that scope, we might set up options. We have server options go out to anyone who's a client of that server, we have scope options that go out to anyone that's a client of that scope--with class options going to anyone that's of a particular class, and client options going to that particular client.

10.1.3 Installing DHCP and Creating a Scope

Installing DHCP and Creating a Scope
0:00-0:04
In this video, we're going to look at installing DHCP and creating a scope.
Install DHCP Server Role
0:05-0:39
The first thing that I need to do is install the DHCP server role. In the Add Roles and Features Wizard, I'm going to hit Next, and I will choose DHCP.
One of the things to notice here is, it tells you right away you should configure at least one static IP address on this computer. The DHCP server must have a static IP address. It cannot be both the server and the client. Then, it's also advising me to plan out what I'm going to do. We'll go ahead and install this.
Complete DHCP Configuration
0:40-0:52
Now, you can see there's a notification up here, and it's prompting me to Complete the DHCP configuration, so that would be your next step.
Security Groups
0:53-1:40
It's going to create security groups for delegation of DHCP server administration. One thing to note about these groups--I've got Administrators and Users. If I install DHCP on a member server, the DHCP Administrators is in the SAM on that member's server.
That means if I add a user to DHCP Administrators, they can run DHCP only on that member server. If it's a domain controller, DHCP Administrators is replicated to all the domain controllers. If I add a user to that DHCP Administrator's group, they are administrators of all the domain controllers that also have DHCP. If you're looking to limit DHCP administration to one server, that server has got to be a member server, not a domain controller.
Authorize DHCP Server
1:41-3:03
This wizard is going to give me the ability to authorize the DHCP server. It says, "if domain joined". Well, it needs to be joined to a domain if there is Active Directory, and it needs to be authorized if there is Active Directory. If there's Active Directory in the environment and it's not joined to a domain and authorized, it's going to shut down. Microsoft does that so that you cannot use windows server to create a rogue DHCP server, which is a DHCP server that's not approved by network administration.
I'm going to go ahead and hit Next, but I will authorize the server later. You can see, I'm not logged in as somebody who has the right to authorize it. To authorize the DHCP server, you must be logged into the domain as a member of the Domain Admin's group, and that you should know. In my case, I've purposely logged on as local user so I can show you how to authorize the server as a separate step.
If I install it as a domain admin, and particularly if it's on a domain controller, it's going to be automatically authorized, and I wouldn't be able to show that to you in the software. We're going to skip it.
We've created our security groups. I'm going to hit Close. It does prompt me to restart the DHCP server service, and I will do that as soon as we get into DHCP administration.
Authorize DHCP Server as a Local User
3:04-4:23
I'm going to open the DHCP counsel using the Tools menu, but you could also use the Start menu. What you want to observe is that both IPv4 and IPv6 have these red arrows that point downwards. That means that the server is not authorized. You should be familiar with what that icon indicates to you.
If I right click my server, I do not have the option to authorize it, because I'm not logged in with somebody who has those rights. I'm logged in as a local account. To bypass this, I'm actually going to do a Run-as and reopen my DHCP counsel using domain credentials.
What I'm going to do is, I'm going to right click DHCP, then I'm going to come down here and Run as a different user. Now, if I right click my server, I can authorize it. You're going to have to hit F5, and once you've got check marks, then you know you're good to go.
Problem: Icons Remain Unauthorized
4:24-5:10
In rare circumstances--it's very irregular and in each generation of Windows Server it happens less--but in rare circumstances, you hit Refresh, the arrows stay red pointing down, but you right click, and it says Unauthorized. You think, well, it should be Authorized, but the arrows didn't change. The icons didn't change. If the icons don't change, it's not going to hand out IP addresses. You could reboot the server, but hopefully any students of mine would say

it's bad to reboot servers, because we don't know if they're going to come back. A much easier way would be to restart the DHCP Services. Even if the icons had not changed, they would change after restarting the services.

Create a Scope

5:11-5:30

Once we have installed DHCP, created the user groups, and authorized the server, our next step is to create a scope. I'm going to right click IPv4 and do a New Scope. I'll hit Next. Now I'm going to put in an IP address range.

IP Address Range

5:31-5:51

This particular computer is on the 192.168.1.0 network. There's only one scope per subnet, and I'm going to allow my scope to run from 1 to 254. I'm going to give out the entire network.

Subnet Mask

5:52-6:22

Once I identify an address range and a subnet mask, and I finish this wizard, I'm not going to be able to change the subnet mask. I can always change the range, but the subnet mask applied to that range gives me the network ID, and once that's been identified, you can't change that after the fact. You'd have to delete and recreate the scope. Only one scope per network, that's why they don't let you change it.

I'm not going to Add any Exclusions or Delay, so I'm simply going to click Next. Each scope has a lease duration, and the lease duration is how long clients can use that IP address before it expires.

Lease Duration

6:23-7:29

The clients will attempt to renew at 50%, and I think it's 87.5% of that lease, and then again at the end; otherwise, it will just go back and go through the process of getting a new IP address.

The politics of the lease duration are this: If you have a very dynamic environment with a lot of users moving around, you want to keep that lease short. Otherwise, the DHCP server may hand out a lot of IP addresses, the clients leave, and eventually the scope runs out of IP addresses. You know you're out of IP addresses because instead of the checkmark icon that you see there, you would see a blue circle with a white exclamation mark that generally means "out of addresses".

If I have a stable environment, I want a longer lease, because why should I have all this DHCP broadcast traffic on the network frequently when the clients hardly ever move. So a stable environment, long lease, dynamic environment, lots of laptops, short lease.

Configure DHCP Options

7:30-7:36

Now I have the option to Configure DHCP Options. I can say or I can configure them later. We'll just finish the wizard.

Default Gateway

7:37-7:43

I can set up a default gateway for this particular network and add that in. That would be my router.

DNS Servers

7:44-7:58

I can also set up DNS servers, and you can see it's already grabbed the settings from this computer to decide what the parent domain would be and the DNS server should be for this particular network, but I could change it if I need to.

WIN Servers

7:59-8:59

I have the ability to set a Win Server if I'm using Wins in my environment. Hopefully, you're not. At the end, it asks me if I want to activate the scope. I'm going to say no just so I can show you the icons; what it is. Now, you can see my blue circle with an exclamation mark, which means that something's wrong. Most commonly, that's associated with, "out of IP addresses". Well, here I'm out of IP addresses because I have not activated the scope. I'm going to right click, and I'm going to activate the scope, which means I can actually give out IP addresses from that scope. My address pool is configured using the wizard. I can actually change the addresses in here. I can change the range, but you can see that I cannot change the subnet mask. I can change the lease so that I can change the actual range of IP addresses. I can actually have DHCP register DNS information with the DNS server on behalf of the clients.

Network Access Protection

9:00-9:27

I can integrate the scope with Network Access Protection. Network Access Protection looks at the health of the computer. If the computer is healthy, it gets an IP address. If it's not, it doesn't. Then I can go in and talk about what clients I will service.

Any additional information beyond an IP address and a subnet mask is done as an option. We have four different layers of options.

Server Options

9:28-9:32

Server options apply to everyone who is a client of that server.

Scope Options

9:33-9:53

Scope options apply only to clients of that scope. You can see that the wizard has actually gone through and set up DNS servers and the DNS domain name as being a scope option. Many times, those are server options. They certainly could be scope options. Router is always a scope option because it's unique to that network.

Configure Options,

9:54-10:56

Here, if I want to Configure Options, I can right click Configure Options, and maybe I say, well, DNS Servers and the DNS Domain Name, those should not be Scope Options, because I have multiple networks, but everybody should be pointing to the same DNS server. I could do that as a Server Option which applies to anybody who gets an IP address from my DHCP server. Again, we just right click, Configure Options; and then I would add my domain name. You can see its a little bit different icon. If I click on Scope Options, I will see my Server Options in there, and the difference is, Server Options go to all clients of the server. Scope Options go to just clients of the scope.

Class Options

10:57-11:10

There are also Class Options that you can do, which only apply to clients of a particular class, and then you can set Client Options that only apply to that particular client, which would be done inside of a reservation.

Policies

11:11-12:13

New with Windows Server 2012, we also have Policies. If I had more complex criteria for giving out options, I could make a policy and I can Add conditions. I could say, well, I'm only looking for if the User Class is Network Access Protection, or Default BOOTP, and if either of those things are true, I can set a specific option in that example. If you have some type of a logic equation that governs whether or not they get the option or not, then you could do that with a policy. There are a couple of options that I do want to show you. Let me go back into options.

Microsoft would like you to be aware of the options to support PXE Boot clients.

Options to Support PXE Boot Clients

12:14-13:48

PXE Boot is spelled PXE. PXE Boot clients are clients that can boot off their network card and connect up to a server and download an operating system, either because it's a diskless work station, meaning it does not have a hard drive and it needs to download the operating system because that's its only operating system, or because its going to be used to image the computer for the first time, so it's going to connect up to an imaging server and pull down the operating system that will be installed on the hard drive. Microsoft's role that does that is called Windows Deployment Services: WDS.

To support PXE Boot clients, there are three different options. One of them is not here. you would have to manually add it, but 066 would be the name of the WDS server, whatever that server would be, and then the boot file name would be the name of the file that it should actually pull down to go to the next step in that process. It would be relative to a folder up on the client; something along the lines of that. There is another option, 060, but that's only used if WDS is running on the DHCP server. Again, for Microsoft trivia, 066 and 067 support PXE Boot clients. For real life, the WDS server will go out and program the PXE Boot options, and you don't need to do it manually. That's how we get DHCP installed.

Summary

13:49-14:21

Make sure we authorize the server in Active Directory, otherwise, its going to shut down, and if you have a DHCP server that refuses to give out addresses, shut down, or one that's got those little red arrows, that's your problem. You're not authorized. If there's Active Directory and it's a standalone server, that's your problem. You don't belong to the domain. You're not authorized. Then we create our scope, one scope per subnet. If you want to give out extra information, that's done using Options. That's how we initially get DHCP up and running.

10.1.4 DHCP Facts

Dynamic Host Configuration Protocol (DHCP) centralizes management of IP address assignment by allowing a server to dynamically assign IP addresses to clients. DHCP also allows users who move from network to network to easily obtain an IP address appropriate for their network connection.

The DHCP server and the client use broadcasts to communicate. The table below describes the method clients use to obtain an address from a DHCP server.

Broadcast	Description
DHCP Discover (D)	The client begins by sending out a DHCP Discover frame to identify DHCP servers on the network.
DHCP Offer (O)	A DHCP server that receives a Discover advertisement from a client responds with a DHCP Offer. The offer contains the IP address. If more than one DHCP server sends an offer packet, the client responds to the first offer packet that it receives.
DHCP Request (R)	The client accepts the offered address by sending a DHCP Request.
DHCP ACK (A)	The DHCP responds to the request by sending a DHCP ACK (acknowledgement).

If the DHCP server is across a router, additional implementation steps are required.

The following table identifies DHCP authorization requirements and the authorization verification process.

DHCP Authorization	Description
Authorization requirements	Authorization requirements for a DHCP server include: • Authorization is required if you are using Active Directory; no authorization is required for a standalone server. • When using Active Directory, DHCP servers must either be domain controllers or domain member servers to authorize them for DHCP. • When authorizing a DHCP server, its IP address is added to a list of authorized DHCP servers maintained in Active Directory.

	To authorize a DHCP server, you must be logged in as a member of the Enterprise Admins group. If you install a DHCP server as an Enterprise Admin, the server is automatically authorized.
Authorization verification	Keep in mind the following about DHCP server authorization verification: • When a DHCP server starts, its IP address is compared to the Active Directory list. If it is found, the server is allowed to issue IP addresses. If it is not found, the server automatically shuts down before completing the startup process. • A Windows DHCP server checks for authorization when it boots and reauthorizes every five minutes. • DHCP servers running other operating systems (for example, Unix, NetWare, or Windows NT) do not check for authorization before assigning addresses. • You can authorize a server before or after DHCP is installed.

The DHCP Server role must be installed to set up a DHCP server in Active Directory.

Keep in mind the following when configuring a DHCP Server.

- Configure the DHCP service to autostart.
- The DHCP Server must have a static IP address.
- When you set up DHCP on a member server and add a user to the DHCP Administrators group, that user has DHCP Administrator rights only on the member server. If you delegate administration on a domain controller, the DHCP Administrator has rights on all DHCP servers in the domain.
- Static IP addresses are recommended for DNS servers and domain controllers.

To configure a DHCP server to deliver IP addresses, you must configure the scope. A *scope* is the range of IP addresses that the DHCP server can assign to clients. Be aware of the following when working with scopes:

- There is only one scope per network segment.
- The scope must be activated before the DHCP server will assign addresses to clients. After you activate a scope, do not change its range of IP addresses.
- A scope has a subnet mask that determines the subnet for a given IP address. You cannot change the subnet mask in an existing DHCP scope. To change the subnet mask used by a scope, you must delete and recreate the scope.
- TCP/IP configuration parameters are designated in options.
- Lease duration values are part of the scope properties and determine the length of time a client can use the IP address leased through DHCP.

In addition to providing an IP address, the DHCP server can also provide clients with additional IP configuration parameters using *options*. Commonly used DHCP options include the subnet mask, the default gateway address, and a DNS server address. There are four levels of options that can be configured:

- **Server options** are applied to all computers that get an IP address from the DHCP server, regardless of which scope they obtain the address from. For example, if your organization has only one DNS server, then all DHCP clients need the same DNS server address. This can be done most efficiently with a server option.

- **Scope options** are applied to all computers that get an IP address from a particular scope on the DHCP server. For example, because scopes are associated with specific subnets, each scope needs to be configured with the appropriate a default gateway address option.
- **Class options** are applied to all computers that are members of a particular class. To do this, the class must first be configured individually on each computer so it knows what class it belongs to. Class options are not commonly implemented.
- **Client options** are applied to a specific DHCP client. The client's MAC address is used to identify which system receives the option.

The DHCP console provides context-sensitive icons to reflect DHCP server status as follows:

- A check mark in a green circle indicates the DHCP server is connected and authorized.
- A red down arrow indicates the DHCP server is connected but not authorized.
- A horizontal white line inside a red circle indicates the DHCP server is connected, but the current user does not have the administrative credentials necessary to manage the server.
- An exclamation sign inside a yellow triangle indicates that 90% of available addresses for server scopes are either in use or leased.
- An exclamation sign inside a blue circle indicates 100% of available addresses for server scopes are either in use or leased.

10.2 DHCP Exclusions and Reservations

As you study this section, answer the following questions:

- How are *reservations* different from *exclusions*?
- What are the two ways to exclude IP addresses from a scope?
- What information is necessary to configure a client reservation?
- How are filters used to identify which computers can be clients of the DHCP server?

After finishing this section, you should be able to complete the following tasks:

- Configure exclusion ranges and client reservations.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  Manage DHCP Scopes, Exclusions, and Reservations
  - Create and Configure Scopes
  - Create Exclusion Ranges
  - Create Client Reservations

This section covers the following 70-410 exam objective:

- 402. Deploy and configure Dynamic Host Configuration Protocol (DHCP) service.
  This objective may include but is not limited to:
  Create and configure scopes
  Configure a DHCP reservation
  Configure DHCP options
  Authorize DHCP server

10.2.1 DHCP Exclusions, Reservations, and Filters

DHCP Exclusions, Reservations, and Filters
0:00-0:05
In this video, we are going to talk about exclusions, reservations, and filters.

Exclusions
0:06-1:02
When you create a scope, you can only have one scope per subnet. That makes it a little bit of a problem if the boss comes in and says, "I don't want you to give out all those IP addresses. I want some of them to be reserved for static IP addresses." In that case, we can make an exclusion.
Basically, it would look something like this. When we set up our scope, let's say I am working with 192.168.1.0 network /24 and the boss comes in and says, I want you to hand out addresses 20-100 and 150-200, but if you hand out 1-19 or 101-149, you're fired. What I would do is create a scope that runs from 1-200, but then I would exclude 1-19 and 101-149. Exclusions are IP addresses that are not handed out.
In reality, probably the boss isn't going to come in and say, "you can't hand out these IP addresses". What do we really use them for in real life?

When are Exclusions Used
1:03-2:04
Exclusions are used when we have static IP addresses in the environment. Some servers have to have a static IP address. For example, the DHCP server itself has to have a static IP address. It can't be its own client; that wouldn't work. We also recommend static IP addresses for DNS servers...domain controllers. The advantage of a static IP address is, it's not dependent on DHCP. If my DHCP server goes down, any server that has a static IP address won't have any problems. The disadvantages are, if anything changes in my environment, I'm going to have to visit that server and make the change manually.
What I like to do is create a scope that runs the entire network, input in exclusions or the addresses that I have handed out statically. What this does is create excellent immediate documentation. If I win the lottery and don't come to work tomorrow, the person that gets hired to take my place can immediately look at the scope and see addresses 1-200 are in play in this environment, but 1-19 and 101-149 have been handed out as static addresses. That's what exclusion does for us.

Reservations
2:05-2:58
Let's look at reservations. Reservation is pretty much just like the kind of reservation you would make for a hotel room. If I call up and reserve a hotel room, it guarantees me a room. In the case of IP addresses, a reservation guarantees a particular IP address for a particular client.
How does the DHCP know that particular client is the one requesting the IP address? Every network card has a built-in address called the MAC address that's unique. When I make my reservation, I'm going to go out and get the MAC address of that particular client and put it in the reservation. The other half of the reservation is some IP address in that scope that will be given to that client only.
A couple of things to take note of: you do not make an exclusion for an address that you use in a reservation. Exclusion trumps a reservation. Exclusion means it will not be handed out, not to anyone. Reservation means that particular IP address will only be given to that particular client.

When are Reservations Used
2:59-3:20
Why would I use a reservation? There are some devices or computers that I want to always have the same IP address, but I'm not interested in setting a static IP address. For example, printers. You probably don't want to visit the printer and set a static IP address, but I need to guarantee that that IP address doesn't change, because if it does, I'm going to have to reconfigure all my client computers. That would be a great situation to use a reservation.

Filters
3:21-3:35
The last thing we are going to talk about in this video are filters. Filters were new with Windows Server 2008 R2. Basically, what they do is allow us to control who can be a client of the DHCP server. There are two ways to use filters.

Allow Filters
3:36-3:51
We can go in and use Allow Filters, which then means that only the people on my Allow list will be allowed to be clients of the server. What I would need to do once I turn that on, is to go out and get the MAC address of every single computer that's going to be a client of that server inputted into the filter.

Deny Filters

3:52-4:16

In addition to Allow Filters, we also have Deny Filters. They work the same way. Anybody who is on the Deny list will not be able to be a client of the server.

For some reason I find out that there is a computer. It's obtaining an IP address for my address for my DHCP server, not one of my computers. I don't know who it is, but I just don't want to give an address to that MAC address. I can turn on Deny Filters, put in that MAC address on the Deny list, and those clients will not be allowed to get an IP address.

Summary

4:17-4:42

Exclusions are for computers that we have static addresses, these are computers that should not be dependent on DHCP. They need their IP address no matter what. You put the exclusion in the scope range, just so that we document the fact that they're static IP addresses. Reservations, on the other hand, guarantee a DHCP client a particular IP address. These are for devices that should be DHCP clients that always need the same IP address. Finally, filters allow me to control who is allowed or not allowed to get an IP address from my server.

10.2.2 Creating a DHCP Exclusion, Reservation, and Filter

Creating a DHCP Exclusion, Reservation, and Filter
0:00-0:05
In this video, we're going to look at DHCP exclusions, reservations, and filters.
DHCP Console
0:06-0:30
First, we need to open up the DHCP console. I'm going to do that through the Tools menu, but you can also use the Start menu. You can see that we have a scope here for the 192.168.1.0 network, only one scope per subnet. I currently have an Address Pool of 192.168.1.1 through 254.
DHCP Exclusions
0:31-1:02
DHCP exclusions are used to show or document static IP addresses in the network. It's possible to go through if all your static IP addresses are at the beginning of the range--let's do 1 through 50--I could've done a scope that would run from 51 to 254. If I make an exclusion, anybody who's going to be running DHCP can come in here. When they see the exclusion, they'll know right away that those addresses are being used, but they're being used as static IP addresses in the network.
Creating Exclusions
1:03-1:35
Let's create some exclusions. I'll do a new exclusion range. Let's say that my router is 192.168.1.1. If you just add one IP address, it will add a single IP address. I also know that I have some servers running 40 and then 50 and 51, and so now those IP addresses are excluded from distribution. They absolutely will not be handed out under any circumstances.
Static IP Addresses
1:36-1:59
The politics of it are this. A static IP address is great for servers like DNS servers, DHCP servers, domain controllers--servers that should not be dependent on DHCP. The only problem with using static IP addresses is, if anything changes, like DNS server changes, any of that information should change, I'm going to have to remember which computers have static IP addresses and then manually reconfigure them.
Reservations
2:00-3:58
There are other types of devices that should always retain the same IP address, but it's not necessary that they continue running if DHCP fails. For example, printers. I don't want to revisit every printer in the company if something changes, but on the other hand, I have to make sure that they always get the same address from DHCP; otherwise, I may have to end up reconfiguring printers.
Clients that should maintain the fact that they're DHCP clients but always get the same address. That's done using a reservation. How does the DHCP server know that it's this particular computer that's asking? It's done by providing the MAC address.
To get the MAC address, you want to use your ipconfig /all command. Here's my MAC address; it's the physical address. This is a unique number that's burned into the network card. Every network card has a unique number, and that's how DHCP is going to know which client this is. I'm just going to copy it, so I'm going to right click and do a Mark, highlight it, and then I hit Enter. It goes right into the clipboard.
I've got to have that MAC address when I create my reservation. I'm going to right click and do a New Reservation. I identify which IP address this should get, and then I put in my MAC address. You definitely want to know a reservation is a combination of an IP address and a MAC address-- guarantees that that client will always get the same IP address. This IP address now will not be handed out to any other client in this particular MAC address, and that's how I will guarantee it.
Do not exclude reservation addresses. If it's excluded, it won't be handed out at all. If it's a reservation, it will only be handed out to that computer. I'll add my reservation, and then once I have a reservation, I could actually do individual options for just this client. That's how we do our client options.
Filters
3:59-4:09
If we need to control who is allowed to get an IP address from this server, that's done using filters. We have two types of filters; Allow filters and Deny filters.
Allow Filters
4:10-4:39

Allow filters would be tough to support. You'd have to be in a really secure situation. Basically, what that would say is, only those computers that are listed in the allow list would be allowed to be clients of this DHCP server. I would Enable the allow list, and then I would add in. The problem is I would have to get the MAC address of a single computer in my environment and add them all in.
That might become a little difficult to support. Perhaps you have a problem where somebody is picking up an IP address; you don't know where that computer is.

Deny Filters

4:40-5:04
You're convinced it's somebody that's plugged into your network and shouldn't be, then, it might be very easy to set up a deny filter. It's the same exact situation. Once I enable the deny filter, I then create my filter and provide the MAC address of everyone who should be denied.
Exclusions are used for static IP addresses in the environment.

Summary

5:05-5:20
Reservations are used for DHCP clients who need to be guaranteed the same address--combination of a MAC address and an IP address. Then, filters let me control who's allowed to actually be a client of my DHCP server.

10.2.3 DHCP Exclusion and Reservation Facts

DHCP exclusions, reservations, and filters help to control DHCP IP address assignment.

Mechanism	Description
Exclusion	Use *exclusions* to prevent the DHCP server from assigning certain IP addresses. • The scope lists the range of IP addresses for the network. Exclusions identify those IP addresses within the range that are excluded, such as an address used as the static IP address for a server • You set exclusions by right-clicking **Address Pool** and selecting **Exceptions**.
Reservation	Use *reservations* to make sure a client gets the same IP address each time from the DHCP server. For example, use a reservation for printers to keep their IP addresses consistent while still assigning the addresses dynamically. • The reservation associates the MAC address with the IP address the client should receive. • You can use the **ipconfig /all** command to determine the MAC address of a computer. • You can configure client options for the individual computers listed under reservations. When using reservations, do not exclude the addresses you want to assign. Excluded addresses are not assigned.
Filter	Filters control which computers can be clients of the DHCP server. Filters use the MAC address of the computer. The two types of filters are: • **Allow** specifies the computers that can be clients of the DHCP server. • **Deny** specifies the computers that cannot be clients of the DHCP server. Enable the **Allow** or **Deny** filter and then specify the MAC and IP addresses.

10.3 DHCP Centralization

As you study this section, answer the following questions:

- When might you choose to create super scopes?
- What type of a system is multinet when discussing *super scope*?
- What are the steps to create a *split scope*?
- What options are available for a DHCP server to service a subnet separated with a router?

After finishing this section, you should be able to complete the following task:

- Configure a DHCP Relay Agent.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  - Configure DHCP Options
    - Configure Server Options
    - Configure Scope Options
  - Implement DHCP Centralization
    - Configure a DHCP Relay Agent
    - Add a DHCP Server on Another Subnet

This section covers the following 70-410 exam objective:

- 402. Deploy and configure Dynamic Host Configuration Protocol (DHCP) service.
  This objective may include but is not limited to:
  - Configure client and server for PXE boot
  - Configure DHCP relay agent
  - Authorize DHCP server

10.3.1 DHCP Centralization

DHCP Centralization
0:00-0:09
We are going to talk about DHCP centralization, and in there, we'll also get into some DHCP backups, and specifically we want to talk about Relay Agents and split scopes.

The Problem with Broadcasts
0:10-1:29
We know when the client boots up, it sends out a series of broadcasts to locate a DHCP server and obtain an IP address from that server. The problem with this system is that broadcasts are not transmitted across routers. By default, routers block broadcasts. If I do want to go with a centralized DHCP solution, where I have multiple networks, I'm going to have to do something to help out the clients that are on the network where there is no DHCP server.
Here I have just drawn a very simple network. I've got my DHCP server, I have given it address 192.168.1.40. It's on the 192.168.1.0/24 network, that's this network here, and we'll just going to go ahead and call this Network A. I've got my router, this interface on the router has an address on Network A. It's also connected to this network over here, which is the 192.168.2.0 network/24, and this has picked up. 2.1 is probably a static IP address, and to make life easier, we will call this Network B.
When this client boots up and sends out its broadcast to locate the DHCP server, this router is going to receive the broadcast, but it's not going to do anything with it because it says, "a broadcast is intended for everybody on Network B, not intended for Network A". I don't need to do anything, it's not my problem.
We've got to do something to help out this client.

Creating Scopes
1:30-2:33
Here is what we will do. This DHCP server, first of all, is going to need two scopes, one for each network. I'll make a scope for A and for B and to keep it simple, let's say our scopes run from 1-100 for each of these networks. Any clients that boot up on A are going to be very happy. They'll just pull from that scope, and the DHCP server knows that those requests come from Network A, because it comes into the network card with this 1.40 address. That's very important.
I see a lot of beginning students make that mistake, where they forget to give the DHCP server a static IP address, and we know if doesn't have a static IP address, it's going to pick up APIPA.
APIPA is anything that starts with 169.254. Let's say this picked up 30.40. If it received a broadcast on an IP address 169.254.30.40, it would look for a scope that starts with 169.254. It's not going to have one, so it wouldn't respond to the clients. That's one kind of a pitfall you can always watch out for with the DHCP; make sure you give the server an address that matches the scope.

RFC 1542 Compliant Routers
2:34-3:09
Let's talk about what we'll do for people on Network B. Something has got to take the broadcast coming out of this client and send it over to this DHCP server, and there are two ways to do this. The easiest way would be to have the router listen for DHCP packets and then relay them or forward them to the DHCP server. If you want to do that, you have to have a router that's capable of it, and those routers are called RFC 1542 Compliant Routers. It's a fancy term that means you can program the router to listen for DHCP traffic and pass it to the other side. It's done with the IP Helper table.

DHCP Relay Agent
3:10-4:32
The Microsoft answer would be to use a DHCP Relay Agent. This would be a server that we build on Network B, and in this server, we're going to install the Routing and Remote Access Service. This would be a server that we build on Network B, and we will install the part of that service that is the DHCP Relay Agent, and it does exactly that. Any broadcasts that come in, my Relay Agent is going to take that and send a direct packet to the DHCP server. Because this Relay Agent has an IP address from Network B, when my DHCP server gets that request, it's going to respond with an address from the B scope and it's going to send it back to the Relay Agent, and the Relay Agent is going to broadcast it on the network, and the client will receive it.
The moral of this story really is, if you have a centralized DHCP solution-- where you have one DHCP server or it doesn't have to be one, it's just the fact that the DHCP server is not on the same network as the client. In that case, we really have two solutions, an RFC 1542 compliant router or the Microsoft solution of building an RRAS service and installing a DHCP Relay Agent.

That having been said, a situation like I have just drawn is not a great solution, because now I have two networks that are dependent on one server. If that server crashes, everybody on both networks are going to be sad.

Fault Tolerance Example

4:33-5:55

Let's look at instituting some fault tolerance for this and how we can beef up the network, and really make it so that no matter what happens, the clients will be covered.

Here I have drawn just two networks, A and B. I'm not going to put in all the IP addresses, and now I've got two DHCP servers, DHCP1 and DHCP2, and let's say we have configured them correctly for each network. Each DHCP server has a scope that's appropriate for its network, and that scope runs from 1-100. I'm trying to make the math easy on me. Same thing over here, DHCP2 is on Network B and it has a scope for Network B that runs from 1-100, but now you start thinking, if DHCP1 fails, everybody on Network A is going to be sad. Same thing for DHCP2, if that goes down, the clients on Network B are going to be out of luck. You start to think, "Wouldn't it be great if these servers can act as backups for each other?" No problem, you can totally do that.

What you need to do is create a scope on each server for both networks, and you are actually going to create the identical scope, but what we don't want is to have conflicts. We can't really let both DHCP servers hand out all the same addresses. What we're going to do is make a gamble, and the gamble is this: one of these machines goes down, we're going to gamble that at that period in time. Only about 25% of my clients really actually need to contact the DHCP server. The rest of them have addresses, their leases haven't expired, everything's great.

Fault Tolerance: 80/20 Rule

5:56-6:23

We're going to go ahead and set up fault tolerance using what they like to call the 80-20 rule, which is that the primary DHCP server for that network is going to handle 80% of the addresses, the secondary server (the backup server on the other network) will just have 20% of them. That way if the primary server fails, the backup server will have 20% of the total addresses, and it will be able to service those clients in the short time it takes me to get the primary server back up and running.

Exclusions

6:24-6:59

I have the same scopes, and what I'm going to do to make sure that duplicate IP addresses are not handed out is to create exclusions.

On DHCP1, I'm going to exclude 20% of the addresses from scope A. I'm going to exclude addresses 80-100. Over on the backup DHCP server, I'm going to exclude the 80%. Now I'm going to the same thing to provide fault tolerance for Network B. Now that I've got my exclusions set up, I know that these servers will not hand out duplicate IP addresses, but the scopes are synchronized. I have one defined scope for each network and I've set up the scope so that they can back up each other.

DHCP Relay Agent

7:00-7:12

The last step that I've got to do is, put a DHCP Relay Agent on each network, because in the event of the primary server failing, the clients on the other network have to be able to contact the backup DHCP server.

Delay Backup Server

7:13-7:50

What I can also do is build in a little bit of a delay on the backup server, so it delays in responding to the clients. The clients will always choose the first DHCP server that they hear responding to them.

There shouldn't be a situation where the backup server responds first before the primary server, but these are computers. I've seen all kinds of crazy things happen in my career. Students say to me, "Shad, why did that happen?", and the only thing I can tell them is, job security. That's why we have jobs, because crazy things happen. You can build in a little bit of delay there so that the secondary never responds before the primary if you don't want that, but we're going to have to put in the Relay Agent on both networks.

Split Scope

7:51-9:02

Historically, if I wanted to set something like this up, I'd have to do all the math myself. Figure out 80% of the addresses, 20%, create the exclusions. The more work I do, the more chance I'm going to make an error. Believe me, if I'm your network admin, I'm not going to be perfect, nobody is. Microsoft with Windows Server 2008 R2 brought in a really excellent technology. I love this.

It's called the split scope, and basically what I do is, on the primary server, I create the entire scope--let's say 1-100--and then I can just actually right-click that scope and say split scope--identify which server will be my backup. I tell it what percentage of the addresses will be kept on the primary, what percentage on the backup. The default is 80/20, but you could do it any way you want. You could do it 50/50, there's no hard and fast rule it has to be 80/20.

Then it will go ahead and create the backup scope on the backup server with all the appropriate information, completely replicating the original scope. The only catch is that I still have to go into the backup server and activate that scope when the split scope wizard makes it for me. It doesn't turn it on, just to, out of a courtesy, let me decide whether it will be always hanging out there, or whether it's something I'm going to manually activate when my primary goes down.

Summary

9:03-9:29

DHCP centralization--if we have a DHCP server that's on a different subnet, we're going to set up an RFC 1542 compliant router or DHCP Relay Agent on the network where there is no DHCP server, or we can actually go through and have a DHCP server on each network and use a backup scope, so that if the primary fails, the secondary can give out addresses. Just to make it really cool, we can use the split scope wizard, to make that easy, and have the computer program the backup scope for us. All we'll have to do is go in and enable it.

10.3.2 Creating a Split Scope

Creating a Split Scope
0:00-0:09
In this video, we are going to talk about split-scopes and Relay Agents essentially centralizing DHCP and providing some fault tolerance.

Split Scope and Failover Wizard
0:10-0:32
First, we are going to take a look at the split-scope, so we need to go into our DHCP console. I'm going to get it from the Tools menu. You can see, I've got a scope here for the 192.168.1.0 network and my Address Pool runs from 1 to 254. I have some exclusions.

Failover for the Scope
0:33-1:56
There is a new feature with Windows Server 2012 that allows me to configure failover for the scope. I could designate member2.northsim.com as the primary DHCP server. If that server fails, add a backup server that will take over, and that would be done by right clicking and Configure Failover. I would put in my partner server and then I would set up the relationship, so if the backup server should get a little bit of lead time, I can set that up as a delay. I can set up load balancing, which means that both DHCP servers will respond to DHCP requests, but they will synchronize their address leases so that they don't hand out conflicting IP addresses. Or I can set up my backup server as a Standby in which case, it is only going to take over if the primary fails. It could be an Active Standby or it can just be Passive and then I can state how long will it take to switch over.
Then I need to give some kind of a Shared Secret that they will use to enable authentication.

Failover Relationship
1:57-2:24
If I create a failover relationship between these two servers, what it means is they will synchronize their database of address leases, so at any given time they should both know which clients have picked up which IP addresses and when those IP address leases fail and know when those IP addresses expire. I'm not going to do a failover, but I wanted you to see that because it is new with Server 2012.

Split Scope
2:25-3:09
Split-scope is a little bit different. In a split-scope, I'm actually configuring a backup DHCP server but it is not as sophisticated as this failover wizard. What it's going to do is create the same scope on another server, but it is going to distribute the addresses according to a percentage that I specify, the idea being that if the primary fails, the secondary will pick up using another piece of the scope that is not in play. With failover they both have the same scope; they are both handing out all those IP addresses. With the split-scope, they both have the same scope but they are handing out different portions of it, the idea being if the primary goes down, the backup scope will take over.

80/20
3:10-4:10
I can right click here, go into Advanced and choose Split-Scope, welcome to the wizard, Next. I will hit Next and then it asks me how I want to split it. Members serving the host will get 80% of the addresses. DC1 being the backup server will get 20% of the addresses.
What that means is, if member2 has 80% of the addresses, this wizard is going to exclude 20% of them. If you can see it is going to do an exclusion from 204 to 254 which is roughly 20% of the addresses. Because the backup DHCP server is going to have 20% active, it is going to exclude the 80% that is on the primary server, so you can see that is going to create an exclusion running from 1 to 203, and I can adjust that either by adjusting the percentage, the slider, or the actual addresses themselves.
I'm simply going to click Next.

Delays for Backup Server
4:11-4:59
If I'm concerned that the backup server is going to respond quicker than the primary server, I can build in a delay of how many milliseconds I want and maybe 10 milliseconds just to make sure that the clients always contact the host first. Then I will hit Finish and it has actually created that scope on the other DHCP server and configured it for me. If I go into my address pool, you can see that it has excluded the 20% of the addresses. If we look on the other server, we will see that it has done the reverse. Here is the backup scope that has been created by the Split-Scope Wizard and it has got 1 through 203 excluded, so it only has 20% of the addresses available to it.

Activate the Backup Server
5:00-5:44
You should be aware that when you use a Split-Scope Wizard, the backup scope is deactivated by default, so if you want that backup scope to actually do its job, you have to come in and manually Activate the scope. If these were two servers that were backing each other up, and the backup server is located on a different subnet or any time you have centralized DHCP and you have clients that are located on a subnet that does not have a local DHCP server, you have to do something to allow those client broadcast to get to the DHCP server. Either you have a RFC 1542 compliant router or if you have a Microsoft Server, you can set it up as a DHCP Relay Agent.

DHCP Relay Agent
5:45-6:10
DHCP Relay Agent's job is just that, to relay the DHCP broadcast to the DHCP server, so the Relay Agent is going to be configured on the remote subnet where there is no DHCP server. No DHCP server, DHCP Relay Agent.

Routing and Remote Access Role
6:11-6:38
In order to do this, we need the Routing and Remote Access role to be installed. I'm going to Enable Routing. Once we have the role installed, we need to go in and configure it.

Configure Remote Access
6:39-7:22
I'm going to go up under Tools and I want to open up Routing and Remote Access. This tool itself is called Routing and Remote Access. The role I install is just Remote Access. It is little bizarre. The first thing I'm going to have to do is right click it and Configure and Enable Routing and Remote Access. The least amount I can have this to will be click Custom and then check what I want. For this, I'm going to need LAN routing, that's it. I will hit Next. Going to start the service. That's great. Once I have got the service started, I need to add the DHCP Relay Agent in, I'm going to open up IPv4, I'll right click General, and even though it is not a Routing Protocol, we are going to click New Routing Protocol, and that's where we will see DHCP Relay Agent, and that adds that into RS.

Adding DHCP Relay Agent
7:23-7:42
There's two steps to the Relay Agent.

Choosing an Interface
7:43-8:01
One step is to say which network card it is going to listen on. I'm going to right click DHCP Relay Agent and hit New Interface and choose which interface it will listen on. Whatever comes into that particular network card, if it is DHCP traffic, then it will be relayed.

Address of DHCP Server
8:02-8:20
The second step is to tell the Relay Agent where to relay it, so I will go into the Properties of my Relay Agent and give the address of the DHCP server. Now any packets that come in on that interface for DHCP would be relayed by this server to the DHCP server. Two computers actively backing each other up that's failover.

Summary
8:21-8:39
For a centralized DHCP solution, I can use a DHCP Relay Agent for the clients that are not on my network. Then if I have servers that are backing each other up, I can create a split-scope, where they have the same scope but different exclusions, little bit of a delay on the backup, so that it can service clients if the primary fails.

10.3.3 DHCP Centralization Facts

Generally, a single subnet uses only one scope with a single range of IP addresses. However, the following table describes administrative features that allow a DHCP server to support multiple scopes.

Feature	Description
Superscope	A *superscope* allows you to combine multiple network segments, and therefore multiple scopes, into a single logical scope. The superscope combines multiple address ranges into a single, logical range. You might use a superscope to: • Add more IP addresses to an existing scope. Adding addresses using a superscope might be easier than deleting and reconfiguring the existing scope. • Migrate clients from one scope to another scope over a period of time. With the superscope, the old and new scopes coexist for a period of time. • Support clients on a single physical network segment that uses multiple logical IP subnets (this configuration is called a *multinet*). • Place multiple DHCP servers on a single physical network segment, with each server servicing a different logical subnet. • Group scopes for non-local scopes (those serviced by a DHCP relay agent on another subnet). The superscope contains all scopes for remote subnets. The scope can then be activated and managed as a single unit. To create a superscope, open the DHCP snap-in and select the DHCP server on which you want to configure the superscope. Select **New Superscope** from the **Action** menu and follow the instruction in the New Superscope Wizard.

Split scope	A *split scope* (also called a *distributed* scope) allows you to provide fault tolerance and improve DHCP performance. In a split scope, two DHCP servers service each subnet. To create a split scope:
	1. Create a scope on each server with the full range of addresses that can be assigned. This is required to ensure that both servers recognize all valid IP addresses on the network.
	2. Create an exclusion on each server, excluding the range of addresses that the server should not assign.
	  ○ If both servers are on the same subnet, exclude 50% of the addresses on each server to allow each server to answer requests equally.
	  ○ If each DHCP server is on a separate subnet, or to make one server on the subnet the primary DHCP server, exclude addresses using the 80/20 rule. Exclude 20% of the addresses on the preferred server and 80% of the addresses on the backup server. The preferred server will be the server that:
	  ▪ Resides on the local subnet
	  ▪ Is nearest to the subnet (in terms of router hops and connection speed)

	▪ Offers the best response time Make sure that there are enough IP addresses available on the preferred server's scope to service the entire subnet. 3. Configure relay agents or enable BOOTP forwarding to allow DHCP requests from remote subnets to reach the DHCP server when DHCP servers are on different subnets. 4. Activate the scopes. A client computer accepts the first DHCP lease offer it receives. For this reason, you cannot control which DHCP server will be used for the actual assignment. The only way you can control this is to ensure that the preferred server responds before the backup server. When using a relay agent to forward requests, configure a four-second (or longer) delay to give the local server time to respond.

The following table identifies the two options for a DHCP server to service a subnet separated with a router.

Option	Description
1542 compliant router	An RFC 1542-compliant router listens for DHCP traffic and routes received DHCP packets to the appropriate subnet.
DHCP relay agent	A DHCP Relay Agent is installed as part of the Routing and Remote Access service (RRAS) enabled on a server. The DHCP Relay Agent sends DHCP packets it receives to the DHCP server through a router.

New in Windows Server 2012, DHCP failover allows two DHCP servers to support the same subnet or scope. When using DHCP failover:

- DHCP scope information is replicated between the two DHCP servers, allowing one of the servers to take over in case of the failure of the other DHCP server.
- The DHCP failover servers can also be configured for load balancing.
- DHCP failover supports a maximum of two DHCP servers.
- Only IPv4 scopes and subnets are supported by DHCP failover.
- IPv6 network nodes typically use stateless IP auto configuration to determine their own IP addresses. The DHCP server typically provides only DHCP option configuration data.

10.4 DHCP Troubleshooting

As you study this section, answer the following questions:

- A Windows client system in your network has an IP address of 169.254.0.1. Why was this address assigned?
- DHCP discoverers are not increasing the DHCP Console. What could be causing this?
- DHCP discoverers are increasing in the DHCP Console, but all offers are static. What could be causing this?
- DHCP discoverers and offers are increasing in the DHCP Console, but requests are static. What is causing this?

After finishing this section, you should be able to complete the following tasks:

- Configure automatic and alternate addressing.
- Troubleshoot common DHCP issues.

This section covers the following Windows Server Pro: Install and Configure exam objective:

- 7.0 Networking and DHCP.
  Implement DHCP Centralization
  - Configure Automatic and Alternate Addressing

Automatic Private IP Addressing (APIPA)

0:00-0:42

Let's talk about APIPA addresses.

That stands for Automatic Private IP Address. Our DHCP clients when they boot up don't have an IP address, and they send out a series of broadcasts to get an IP address from the DHCP server. DHCP discover. Help, is there a DHCP server in the house?

What happens when no DHCP server responds, either because there isn't a DHCP server, or because the DHCP server has failed? The client is going to try that broadcast a number of times. At some point, it sort of gives up. In that case, it will failover to an APIPA address.

The idea behind APIPA was, if DHCP is broke or not present, at least the computer will get some kind of an IP address that it can use in the interim.

169.254.0.0/16

0:43-1:17

We know it's an APIPA address because it comes from the 169.254.0.0 network with a /16 bit subnet mask. If you have an APIPA, address DHCP broke.

As soon as I do an ipconfig and I see that my IP address starts with 169.254, I know that DHCP isn't working. Essentially, that's the way I think of it. DHCP broke. Then I have to go and figure out why this particular client isn't picking up an address. Maybe the DHCP server isn't broken. There's a problem with the network card or the cable, but it certainly tells me that I have not gotten an address from DHCP and I'm supposed to be a DHCP client.

Alternate Configuration

1:18-2:08

There is one thing I would add into this, which is that you don't have to use APIPA as the alternative if DHCP isn't working. You can go in and set an alternate configuration. Generally, where this would be used is if you have a situation where you have a computer that needs to operate on networks that require DHCP and needs to operate on networks that require static IP address, which is actually fairly rare. In that case, I can set a static IP address that the client will failover to in the event that it can't contact DHCP. APIPA addresses only come with an IP address and a subnet mask. No default gateway, no DNS. Though the computer is not going to be very functional, it'll simply be able to find other computers and local resources in its network, but certainly with an APIPA address you're not going to get very far and you're definitely not going to get out to the internet.

Summary

2:09-2:16

APIPA is a failover; it's an address that we get when DHCP doesn't respond for whatever reason, and we know we have an APIPA address, because it starts with 169.254.

10.4.2 Configuring an Alternative Address

Configuring an Alternate Address
0:00-0:04
In this video, we're going to look at APIPA and configuring an alternate address.
APIPA
0:05-1:21
If your computer is a DHCP client and it cannot pick up an IP address from the DHCP server, you should get an APIPA address. (which starts with 169.254) and the computer will randomly pick the last two digits.
Let's take a look at that. I can either right click down here, go to Network and Sharing Center, or I can go in here and click on the Settings that will bring me to my network adaptor. Right now, this particular computer has a static IP address. If I want to make it a DHCP client, I'll tell it to obtain an IP address automatically. You can see, as soon as I do that, the Alternate Configuration tab gets colored in, and what's selected is Automatic private IP address. We're going to leave that configured that way for a minute, and we'll take a look at the IP address that the client gets. You can see I've got 169.254. I happened to pick up 181.97. It uses the default 16-bit subnet mask, and there's no default gateway, no DNS server. This is literally just an IP address, so I could find other local resources on the network until the network administrator fixes the DHCP server. 169.254, that's APIPA. APIPA means DHCP broke, and I probably need to fix that.
A Client that Needs Both a DHCP and a Static IP Address
1:22-1:48
The alternate configuration can also be used if you have a client that needs to both use DHCP and a static IP address. This would probably not be a server; it would be a client. It falls under DHCP, so it's something that you should be aware of. If I need to be a DHCP client--let's say at home I'm using DHCP with my home router--then I want to leave it on DHCP. If at work, all the computers need a static IP address, I don't want to have to keep running into this box and changing it.
Static Alternate Configuration
1:49-2:57
That's why we have the ability to set a static Alternate Configuration.
We'll go back into Properties. Instead of using APIPA, I can define an address that would be used in the event that the computer needs to failover. We'll leave it set to 60. Those settings come in because I've actually set one before in the past. Otherwise, you'd have to type them in. It won't just pop up. Because it's not contacting DHCP, it is running Network Diagnostics, but I'm going to cancel that because I know nothing's wrong.
Now, even though DHCP Eabled is Yes, I have an address that's not APIPA. If you're trying to troubleshoot screen shots or diagrams, make sure, you know, it tells you if DHCP is enabled or not. If it's enabled Yes, and I've got this address, then I would look for the address of the DHCP server. Since there's no DHCP server listed, I can tell that this is the alternate configuration. If, in fact, this was an IP address from the DHCP server, it would list the address of the DHCP server and when that reservation expires.
Network Connection Details
2:58-3:16
Just by looking at Network Connection Details, you can tell if you're a DHCP client, if you've got an APIPA address, if you get an alternate configuration, or you've successfully obtained an IP address from DHCP, or if it's a static IP address. Those are really the four combinations you could possibly have. That's how we set up a static alternate configuration.

10.4.3 DHCP Troubleshooting

DHCP Troubleshooting
0:00-0:03
In this video we're going to talk about DHCP troubleshooting.
DHCP Console
0:04-0:19
I need to open up the DHCP console. Generally speaking, I can tell everything I need to know about what's wrong with DHCP right from this console. What you do is, right click and Display Statistics.
Display Statistics.
0:20-0:31
The only thing that's annoying is clicking Refresh; it doesn't actually dynamically refresh. I can see all of my statistics here.
Discoverers Not Going Up
0:32-1:07
If DHCP Discovers are not going up, that indicates a hardware problem, because what it means is, my DHCP server is not even getting the broadcast. A hardware problem could be a network card on the server or the client. It could be a network cable, could be a switch, could be a router. If the clients are on another network, it could be the DHCP Relay Agent. But, at a hardware level, I am not getting the request. If the DHCP server doesn't get the request, it is certainly not going to give out an answer. In response to that initial Discovers broadcast, the DHCP server should make an Offer.
Discovers Going Up but Offers Remain the Same
1:08-1:27
If Discoverers are going up, but Offers remain the same, in my case, everything is zero, but let's say Discoverers keep going up, and I've got 10 Offers out, and it stays at 10. What that tells me is, the server is getting the request, but it's not responding. Now, why wouldn't the DHCP server respond? There's really only two reasons.
Our of Addresses
1:28-1:41
One reason would be that it's out of addresses, and I should know that right away, because instead of the green checkmark, I would see a blue circle with a white exclamation mark. In case I didn't notice that, that would be one reason why it might not make an offer.
Doesn't Believe it Can Help the Client
1:42-2:38
The only other common reason is that it doesn't believe it can help the client. Why would it not believe that? Suppose the IP address of this particular server was 192.168.2.40. My scope is for the 192.168.1.0 network. So, a broadcast that comes in to a network card--with a network ID of 192.168.2.0--the server assumes that client is not a client of the 192.168.1.0 network. It doesn't have a scope for the 2.0 network, therefore, it doesn't make an offer. It would be a mismatch between the server's IP address and the scope; or, if the request was coming in from a relay agent, between the relay agent's address and the scope. You need to make sure, if the DHCP server is on that network, that it has a good IP address for that network, it matches the scope. The relay agent has to have an IP address that's going to match the network for which it's relaying.
Discovers and Offers Going Up but Requests Not Going Up
2:39-3:20
If Discoverers are going up, and Offers are going up, but Requests are not going up, the only thing that it could be would be a rogue DHCP server. The client picks the first DHCP server that responds. If my server got the initial request, it made an offer, but the client didn't choose it. It must have chosen something else. In that case, that would be another DHCP server. I'm not aware of one; it's a rogue DHCP server. It can't be a hardware problem, because we've already exchanged a couple of broadcasts. That's pretty much it. You would not see a situation where all three of those were going up, but then Acknowledgements were not going out. That to me, is to restart the DHCP Server Service.
Summary
3:21-3:28
Just a quick look at troubleshooting DHCP; if you go in to statistics, you can get quite a bit of information just out of this one little box.

10.4.4 DHCP Troubleshooting Facts

When DHCP is not working, computers on your network will not be able to communicate. The following table identifies methods by which hosts obtain IP addresses in the event of a DHCP failure.

Method	Description
Automatic Private IPv4 Addressing (APIPA)	APIPA is an automatic configuration method in which hosts automatically select their own IPv4 address. • The client computer uses APIPA if a DHCP server cannot be contacted. • APIPA uses IP addresses in the 169.254.0.1 to 169.254.255.255 range. • APIPA sets only the IPv4 address and subnet mask. Because it does not assign a default gateway, APIPA communication is limited to a single subnet.
Alternate IPv4 configuration	With an alternate IPv4 configuration, the system attempts to use DHCP for TCP/IPv4 configuration information. If a DHCP server cannot be contacted, the static configuration values are used. Alternate configuration is set up in the network adapter properties. When you configure an alternate IPv4 address, APIPA is no longer used.

The DHCP console provides valuable information for troubleshooting DHCP. To use the DHCP console for troubleshooting, choose the **Display statistics** option for the IPv4 object. The following table identifies potential problems based on the statistics displayed. Use the **Refresh** button to update the statistics.

Statistics	Potential Problem
DHCP Discovers are not increasing	There is a hardware failure that is preventing communication with the DHCP server. Potential problems include: • Network adapter failure • Network cable failure • Failed switch • Problem with a router • Problem with a DHCP Relay Agent
If Discovers are increasing, but Offers are static	The server is not responding to requests it receives. Potential problems include:

	• The DHCP server is out of addresses. An out of address condition is indicated by a blue circle with a white exclamation mark over the IPv4 object. • The client or the DHCP Relay Agent is not within the scope of the DHCP server.
If Discovers and Offers are increasing, but Requests are static	The client is accepting the offer of another DHCP server, indicating that there is a rogue DHCP server on the network.

Windows Server Pro: Install and Configure Exam Objectives

Windows Server Pro: Install and Configure

This Windows Server Pro: Install and Configure course covers the following objectives:

#	Objective	Module.Section
1.0	Configure Windows Servers • Navigate Server Interfaces Navigate the Windows Server 2008 R2 User Interface Navigate the Windows Server 2012 User Interface • Configure Server Services Configure Server Services Configure NIC Teaming • Configure Server Storage Configure Server Volumes Configure Fault Tolerant Volumes Create a Mount Point Create and Mount Virtual Hard Disks (VHDs) Create a Storage Pool	1.2 1.3 2.1 2.2 2.3 2.4 2.5 2.6 2.7 2.8
2.0	Hyper-V • Manage Virtual Machines with Hyper-V Manager Create and Manage Virtual Machines Create Virtual Hard Disks (VHDs) Create a Differencing Drives (Parent-Child) Configure Virtual Networks and Settings	3.1 3.2 3.3
3.0	Active Directory • Manage Active Directory Configure Global Catalog Servers Create Organizational Units (OUs) Delegate Administrative Control • Create and Manage User and Computer Accounts Create User Accounts Manage User Accounts Configure User Account Restrictions Create Computer Accounts • Create and Manage Groups Create and Manage Global Groups	4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9

	Create and Manage Distribution Group Change the Group Scope Implement Recommended Group Strategy	
4.0	DNS • Manage DNS Configuration Configure Search Suffixes Configure Forwarders Configure Root Hints • Create DNS Zones Create a Root Zone Create an Active Directory-integrated Zone Convert a Zone to Active Directory-integrated • Create DNS Records Create a Zone and Add Records Create A and CNAME Records Troubleshoot Name Resolution 1	5.1 5.2 5.3 5.4 5.5 5.6
5.0	File and Print Services • Manage NTFS Permissions Configure NTFS Permissions Configure Inherited Permissions • Share Folders and Configure Share Permissions Share Folders Manage Shared Folders Manage Share Caching Configure Share Permissions • Manage Combined NTFS and Share Permissions • Configure Volume Shadow Copy Service (VSS) Enable and Configure Shadow Copies Restore Previous Versions of Files and Folders • Configure Quotas Manage Quota Restrictions Create Quota Entries Configure Quota Limits • Manage Printing Create, Share and Manage a Printer Configure Printer Pooling Restrict Printer Access Deploy Printers with Group Policy	6.1 6.2 6.3 6.4 6.5 7.1 7.2
6.0	Group Policy • Manage Group Policy Objects (GPOs)	8.1 8.2 8.3

	Create and Link a GPO Create a Starter GPO Modify GPO Links • Manage Security Policies Configure Security Options Configure Account Password Policies Enforce User Account Control Configure Audit Policies Configure User Rights Configure Restricted Groups • Configure Application Restriction Policies	8.4 8.5 8.6 8.7 8.9 8.10
7.0	Networking and DHCP • Configure Basic Network Settings Configure IPv4 Settings Configure IPv6 Settings • Configure Network Settings for Multiple Subnets Use Subnetting to split address ranges Use Supernetting to combine address ranges Configure Networking for Multiple Subnets • Installation and Authorization Install the DHCP Role Authorize DHCP Servers • Manage DHCP Scopes, Exclusions, and Reservations Create and Configure Scopes Create Exclusion Ranges Create Client Reservations • Configure DHCP Options Configure Server Options Configure Scope Options • Implement DHCP Centralization Configure a DHCP Relay Agent Add a DHCP Server on Another Subnet Configure Automatic and Alternate Addressing	9.1 9.2 9.3 10.1 10.2 10.3 10.4