



Exam
Manual

Microsoft

Server 2008

Network Infrastructure (70-642)



**Smarter
Training**

This LearnSmart exam manual covers the most important concepts and topics you will encounter on the Server 2008 Network Infrastructure exam (70-642). By studying this manual, you will become familiar with a broad spectrum of must-know exam topics, including:

- Configuring IP Addressing and Services
- Configuring Name Resolution
- Configuring Network Access
- Server 2008 R2 Features
- And more!

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

Microsoft Server 2008 Network Infrastructure (70-642) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC.
Product ID: 012379
Production Date: August 2, 2011

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

Warning and Disclaimer

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

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

Volume, Corporate, and Educational Sales

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

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

International Contact Information

International: +1 (813) 769-0920

United Kingdom: (0) 20 8816 8036

Table of Contents

<i>Abstract</i>	7
<i>What to Know</i>	7
<i>Tips</i>	7
Domain 1: Configure IP Addressing and Services	8
Configuring IPv4 and IPv6 Addressing	8
<i>IPv4 – 32-Bit Addressing</i>	8
<i>IPv4 Address Ranges</i>	8
<i>Subnetting IPv4</i>	9
<i>Supernetting</i>	10
<i>IPv6 – 128-Bit Addressing</i>	10
<i>Multi-Home Devices</i>	12
Configuring Dynamic Host Configuration Protocol (DHCP)	12
<i>Configuring DHCP</i>	13
<i>PXE Boot</i>	18
<i>Relay Agents</i>	19
<i>Scopes</i>	19
<i>Exclusions</i>	19
<i>Authorizing a DHCP Server in Active Directory Domain Services</i>	19
<i>DHCP Reservations</i>	20
<i>Windows Server Hyper-V</i>	20
Configuring Routing	21
<i>Setting Up Static and Persistent Routes</i>	21
<i>Using a Routing Protocol</i>	21
Configuring IPsec	22
<i>To configure IPsec:</i>	23
<i>Security Associations</i>	26
Domain 2: Configuring Name Resolution	27
Configuring a Domain Name System (DNS) Server	27
<i>Root Servers and Root Hints</i>	28
<i>Caching-Only</i>	30
<i>Conditional Forwarding</i>	30
<i>DNS Round Robin</i>	31
<i>Netmask Ordering</i>	32

<i>WINS and Windows Server 2008</i>	32
Configuring DNS Zones	32
<i>Types of Zones</i>	32
<i>Aging and Scavenging</i>	34
<i>Clearing the DNS Server Cache</i>	35
<i>NSLOOKUP</i>	35
<i>Dynamic DNS Updating</i>	35
Types of DNS Records	36
<i>Configuring DNS Replication</i>	37
<i>DNSSEC</i>	38
<i>DNS Socket Pooling</i>	38
<i>DNS Cache-Locking</i>	39
Configuring Name Resolution for Clients	39
<i>The Hosts File</i>	39
<i>Selecting a DNS Server</i>	39
<i>LMHOSTS</i>	40
<i>Clearing the Resolver Cache</i>	41
<i>DNS Server Lists</i>	41
<i>Link-Local Multicast Name Resolution</i>	42
Domain 3: Configuring Network Access	43
<i>Network Address Translation (NAT)</i>	43
Configuring Remote Access	47
<i>Remote Dial-Up Access</i>	48
<i>Virtual Private Networks (VPN)</i>	48
<i>DirectAccess</i>	49
<i>Remote Authentication Protocols</i>	50
<i>DHCP Considerations for Mobile Clients</i>	52
<i>RADIUS and Network Policy Server</i>	52
<i>RADIUS Clients, Server, and Proxies</i>	53
Configure Network Access Protection (NAP)	54
<i>DHCP Enforcement</i>	54
<i>VPN enforcement</i>	55
<i>IPSec with Health Registration Authority Enforcement</i>	55
<i>802.1x Enforcement</i>	55

<i>System Health Agents and System Health Validators</i>	55
Configuring Wireless Access	57
<i>Encryption Technologies for Wireless</i>	57
<i>Wireless Protocols</i>	58
<i>Ad Hoc vs. Infrastructure Mode</i>	59
<i>Group Policy for Wireless</i>	59
Configuring Firewall Settings	60
<i>Firewall Profiles</i>	60
<i>Identify Ports and Protocols</i>	62
<i>Firewall Scope</i>	63
<i>Configuring the Firewall through Group Policy</i>	64
<i>Authorizing Connections to Applications without Built-In Access Control</i>	64
Domain 4: Configuring File and Print Services	65
Configuring a File Server.....	65
<i>File Server Resource Manager</i>	66
<i>Quota Management</i>	67
<i>File Screening Management</i>	67
<i>Security</i>	70
<i>Data Recovery Agents (DRA)</i>	74
<i>Publishing a File Share</i>	74
<i>BranchCache</i>	75
Configuring Distributed File System (DFS)	76
Configuring Volume Shadow Copy Services (VSS).....	80
Configuring Backup and Restore	81
Managing Disk Quotas.....	82
Configuring and Monitoring Print Services	84
<i>Configuring Printers with the Graphical Interface</i>	84
<i>Configuring Printers on Windows Server 2008 Core</i>	85
<i>Monitoring</i>	85
<i>File Server Resource Manager</i>	85
Domain 5: Monitoring and Managing a Network Infrastructure	85
Configuring Windows Server Update Services (WSUS)	85
<i>Creating a Computer Group</i>	92
<i>Assigning Clients through Group Policy</i>	93

<i>Auto-Approval Rules</i>	94
<i>Autonomous vs. Replica WSUS Servers</i>	94
Capturing Performance Data	94
Performance Monitor	94
<i>Data Collector Sets</i>	97
Monitoring Event Logs	98
Gathering Network Data	99
<i>SNMP</i>	99
<i>Network Monitor</i>	100
Practice Questions	102
Answers & Explanations	110

Abstract

This LearnSmart Exam Manual was developed to provide you with every procedure, process, and skill necessary to not only take and pass your Server 2008 Network Infrastructure exam (70-642) with confidence, but to adequately prepare you to design and implement the network infrastructure of a Server 2008 system.

We've covered each objective listed by Microsoft, including objectives related to Server 2008 R2, while ensuring that only the most exam-pertinent information is addressed. Inside, you'll find detailed instructions on how to configure IP addressing for your server infrastructure, including planning for IPv6. Additionally, the guide discusses DHCP, routing, IPSec, planning for Name resolution, and Network Access. We will also detail the most important aspect of network administration in Microsoft environment: monitoring and maintaining the infrastructure through systems like WSUS and performance monitoring.

As mentioned previously, we've added material related to Server 2008 R2. With respect to the R2 update, you will find additional information in sections covering IPSec; DNS "Socket Pooling;" DNS Cache-Locking; Remote Access; Virtual Private Networks (VPN); DirectAccess; System Health Agents and System Health Validators; File Server Resource Manager; Branch Cache; Distributed File Systems; and Backup and Restore.

What to Know

Before taking the exam, we recommend having at least a year's practical experience working with Windows Server, preferably Server 2003 or newer. While it may not be possible for you to have on-the-job experience with Server 2008, having professional experience with at least Server 2003 will give you the background knowledge you need to understand the core concepts behind technologies like DHCP and Network Access Protection, for example. Additionally, if your company or organization does not use Server 2008, we would advise you procure a 180-day trial copy of Server 2008 and set up a simple home lab on which to practice some of the configuration routines discussed in this guide. You should also be familiar with Server 2008 R2.

Tips

Obviously, reading and understanding the concepts discussed in this guide is an integral part of preparing for the exam. As mentioned in the **What to Know** section, you should absolutely perform these concepts on a home lab. The key to success on any exam is practice. As such, we would recommend using a practice exam from a reputable vendor, such as LearnSmart. Also, because the manual is designed to deliver exam content and nothing but exam content, it may be worthwhile to take advantage of a more comprehensive study aid, such as video training that will go into extensive detail on the concepts and technologies discussed in this guide. Doing so insures you have a solid background on the content in this manual.

Domain 1: Configure IP Addressing and Services

In the 70-642 Exam, being able to configure IP addresses for a given topology has a large effect on your overall exam score. To do well, you need to be familiar with the concepts of subnetting, supernetting, addressing, IPv4, and IPv6.

Configuring IPv4 and IPv6 Addressing

IPv4 – 32-Bit Addressing

Currently, most IP addresses in use are the 32-bit, IPv4 addresses. There are three different types of IPv4 addresses.

1. **Static Addresses** – These are addresses that have been manually assigned. A static address will remain the same every time that a computer is booted.
2. **Dynamic Addresses** – These are addresses that have been assigned by a Dynamic Host Configuration Protocol, or DHCP, server. With dynamic addressing, it's possible that a computer could have a different IP address every time it gets booted.
3. **APIPA Addresses** – A Windows host can dynamically assign an IP address to itself if a DHCP server is unavailable and a static address has not been assigned. These addresses range from 169.254.0.1 to 169.254.255.254. APIPA addressing can be used for either temporary or ad hoc networks. APIPA addresses are not routable, and hosts that have them can only communicate with other hosts on the same subnet.

IPv4 Address Ranges

In the classful IPv4 address system, there are five classes of IPv4 addresses. (At this time, we only need to be concerned with the first three classes.) The first portion of an IPv4 address identifies the network address, and the second portion identifies the host address. By having more bits dedicated to the host portion, Class A addresses can support more hosts than the other two classes.

Address Class	Number of Network Bits	Available Host Bits	Maximum Hosts
Class A	8	24	16,777,214
Class B	16	16	65,534
Class C	24	8	254

Figure 1: IPv4 Address Class Specifications

Each IPv4 address class has been assigned a certain range of addresses:

- **Class A** 0.0.0.0 to 126.255.255.255
- **Class B** 128.0.0.0 to 191.255.255.255
- **Class C** 192.0.0.0 to 223.255.255.255

IPv4 addresses come in both public and private varieties. Public addresses can be routed across the Internet, while private addresses cannot be routed across the Internet. The following address ranges comprise the private addresses:

- **Class A** 10.0.0.0 to 10.255.255.255
- **Class B** 172.16.0.0 to 172.31.255.255
- **Class C** 192.168.0.0 to 192.168.255.255

Subnetting IPv4

With subnetting, you'll move the boundary between the network and host portions of the IP address to the right. This will create smaller subnets out of a larger classful network address space. The network will then be separated into different collision domains, improving overall network efficiency. Also, if you only need a certain number of hosts on a given network segment, subnetting a larger address space will help prevent wasting IP addresses.

The Subnetting Process:

1. Determine how many host addresses you need per subnet.
2. Find a power of two that results in a number that is greater than the number of required host addresses. This will represent the number of bits that will be used for the host portion of the IP address. For example, if you need 28 host addresses, then you would use 2^5 , which gives a result of 32 addresses. This also means that the last five bits of the IP address represent the hosts. Keep in mind that the first and last addresses of a given subnet cannot be assigned to a host. The first address will be the address of the subnet itself, and the last address will be the broadcast address for that subnet. Therefore, in this example, only 30 addresses are available for assignment to hosts.

The final formula for figuring out how many hosts can be on a given subnet is $(2^n) - 2$.

3. The network address of the first subnet will be the same as the classful network address. Use the same power of two that you used in the previous step to calculate the network address of the next subnet. Add the same power of two to that address to get the network address of the next subnet. Keep doing this until you reach the last possible network address.

To illustrate, let's subnet the 192.168.0.0 network into equal subnets. Our requirement is to have 28 host addresses in each subnet.

Find the power of two that comes closest to your needs. In this case, it would be 25, which equals 32. Subtract two to account for the network and broadcast addresses. That gives us a total of 30 possible host addresses.

Take the 192.168.0.0 address and add 32 to get the next subnet address. Repeat this process until you have all possible subnet addresses.

```
192.168.0.0
192.168.0.32
192.168.0.64
192.168.0.96
192.168.0.128
192.168.0.160
192.168.0.192
192.168.0.224
```

In this case, 192.168.0.224 is the last possible subnet address.

Next we find the broadcast addresses. These are the last host addresses for each subnet.

```
192.168.0.31
192.168.0.63
192.168.0.95
192.168.0.127
192.168.0.159
192.168.0.191
192.168.0.223
```

Our host addresses are everything not reserved by subnet addresses or broadcast addresses.

Subnet .32: 192.168.0.33 - 192.168.0.62
Subnet .64: 192.168.0.65 - 192.168.0.94
Subnet .96: 192.168.0.97 - 192.168.0.126
Subnet .128: 192.168.0.129 - 192.168.0.158
Subnet .160: 192.168.0.161 - 192.168.0.190
Subnet .192: 192.168.0.193 - 192.168.0.222

Supernetting

A supernet is a collection of subnets that are routed by a given router. The network portion of each member subnet address must be the same. Supernets can be used to increase router efficiency, by reducing the overhead involved in routing traffic between different subnets. Supernetting is also referred to as either "route aggregation" or "route summarization."

Let's say that your enterprise is using the 192.168.0.0 address space for its LAN. You've divided this address space into six different subnets, as follows:

192.168.0.0
192.168.1.0
192.168.2.0
192.168.3.0
192.168.4.0
192.168.6.0

In this case, we're using a 16-bit netmask to perform our subnetting. Therefore, the first 16 bits of the address are in common with all of the subnets. Our summarized route, then, is:

192.168.0.0/16

Without supernetting, each router on the network would have to keep track of a very large number of routes in its routing table. With supernetting, one central router can keep track of all the routes, reducing the overhead on all of the other routers.

IPv6 – 128-Bit Addressing

IPv6 addresses are 128 bits long. Sixty-four bits are for the network portion of the address, and 64 bits are for the host portion of the address.

IPv6 Shorthand Addressing

IP version 6 addresses can often be shortened by using a special kind of shorthand. If one 16-bit group consists of nothing but zeros, you can replace it with only a single zero. Also, if consecutive 16-bit groups of an address consist of nothing but zeroes, you can replace those groups with a double-colon. Therefore, you could take an address such as this one:

2001:cdba:0000:0000:0000:0000:3257:9652

and shrink it down considerably.

With shorthand, you can shrink each section of four zeros down to “:0” and then multiple sections of zeros to “::” like so:

```
2001:cdba:0:0:0:0:3257:9652
2001:cdba::3257:9652
```

Note that when two sets of consecutive groups of zeros are separated by a non-zero group, only the first set can be replaced by double colons. For example:

```
2001:cdba:0:0:0:fc34:0:0:9652
2001:cdba::fc34:0:0:9652
```

IPv6 Address Scopes

With IPv6, network packets can be sent to unicast, multicast, or anycast addresses.

Unicast addresses direct network traffic to only a single host. Here are the different types of unicast addresses:

- **Global addresses** - Are routable across the Internet, and operate much the same as IPv4 public addresses. They are reachable on the IPv6 portion of the Internet. The first block of a global address will be in the range of 2000 through 3fff.
- **Unique-Local addresses** - Are the equivalent of private addresses in IPv4. Although these addresses cannot be routed across the Internet, they can be routed across different subnets within a LAN. Unique-Local addresses always begin with “fd”.
- **Link-Local** - Link-local addresses work somewhat the same as APIPA addresses in IPv4. A computer will assign a link-local address to itself. Link-local addresses are non-routable, can only be used on the local subnet, and will always begin with “fe80”. (Note that network adapters will automatically be assigned a link-local address.)
- **Node-Local** - This is the loop-back address for a local computer. It’s always represented by “::1”.

Multicast - addresses direct network traffic to a group of hosts. Sending network traffic to a multicast address requires less bandwidth than sending traffic to each individual unicast address. Multicast addresses can be easily identified because they start with “ff”.

Anycast - addresses are assigned to a specific router function, rather than to a specific network interface. Anycast addresses are for entire groups of hosts. However, traffic sent to an anycast address will only be delivered to the individual host that is closest to the source.

Transitional Techniques

During the current transitional period, there are several available ways to make IPv4 and IPv6 correctly operate with each other.

- **Teredo** - Teredo allows devices with IPv6 addresses to operate across an IPv4 network. It does this by encapsulating IPv6 packets within IPv4 UDP packets. It can be used with Network Address Translation devices, even if those devices cannot be upgraded to IPv6.
- **6to4** - This technique encapsulates IPv6 network packets within an IPv4 shell.

- **ISATAP** – The Intra-Site Automatic Tunnel Addressing Protocol, or ISATAP, creates an IPv6 address stack on top of an IPv4 stack. This allows both IPv4 and IPv6 to access a network device. ISATAP uses a virtual, non-broadcast, multiple-access data link layer that allows it to support multicasting. ISATAP works by taking an interface identifier and preceding it with any valid IPv6 address prefix. An interface identifier can be one of the following:
 - ▶ `::0:5EFE:w.x.y.z` -- In this case, the “w.x.y.z” is a private unicast IPv4 address.
 - ▶ `::200:5EFE:w.x.y.z` -- In this case, the “w.x.y.z” is a public unicast IPv4 address.

In either case, the appropriate prefix still needs to be added. A full IPv6 ISATAP address would look something like one of the following:

```
FE80::5EFE:10.40.1.29
FE80::200:5EFE:64.32.24.200
```

Multi-Home Devices

The term “multi-home” can have different meanings, depending upon the context in which it is used. For our present purposes, a multi-home device is simply one to which more than one IP address has been assigned. This can be done either by assigning multiple IP addresses to a single network interface, or by installing more than one network interface in a given host.

Configuring Dynamic Host Configuration Protocol (DHCP)

The DHCP process is the flow of automatically assigning addresses from a DHCP server to a DHCP client. The process works as illustrated below:

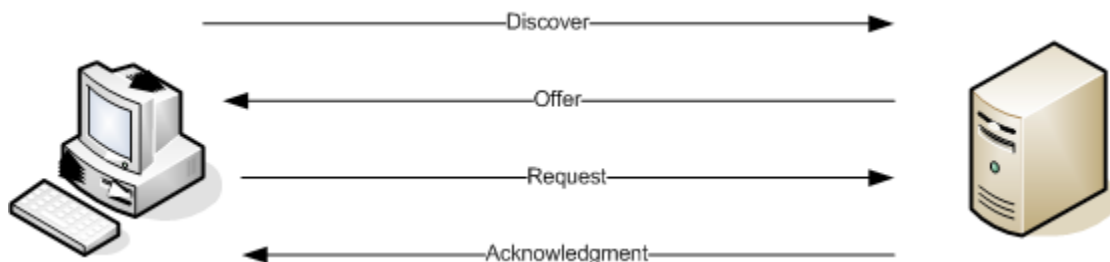


Figure 2: The DHCP Process

The Dynamic Host Configuration Protocol, or DHCP, is used to automatically assign IP addresses to network hosts. When a host is first brought online, it will broadcast a discovery message to find an available DHCP server. Any DHCP servers on the network will then offer an IP address lease to the host. The host will then broadcast a request message, which will let all DHCP servers know which address offer has accepted. The selected DHCP server will then send an acknowledgment message to the host.

Configuring DHCP

With a Graphical Interface

Configuring DHCP on Windows Server 2008 is an easy and straightforward process using Server Manager. To configure DHCP, perform the following steps:

1. Open the **Server Manager**.
2. Select **Roles** and then select **Add Roles**.
3. Choose **DHCP Server** in the **Server Roles Manager** and click **Next**.
4. After reading through the introduction page, click **Next**.
5. On the address you see below, make *sure* that the address is the address of your server.

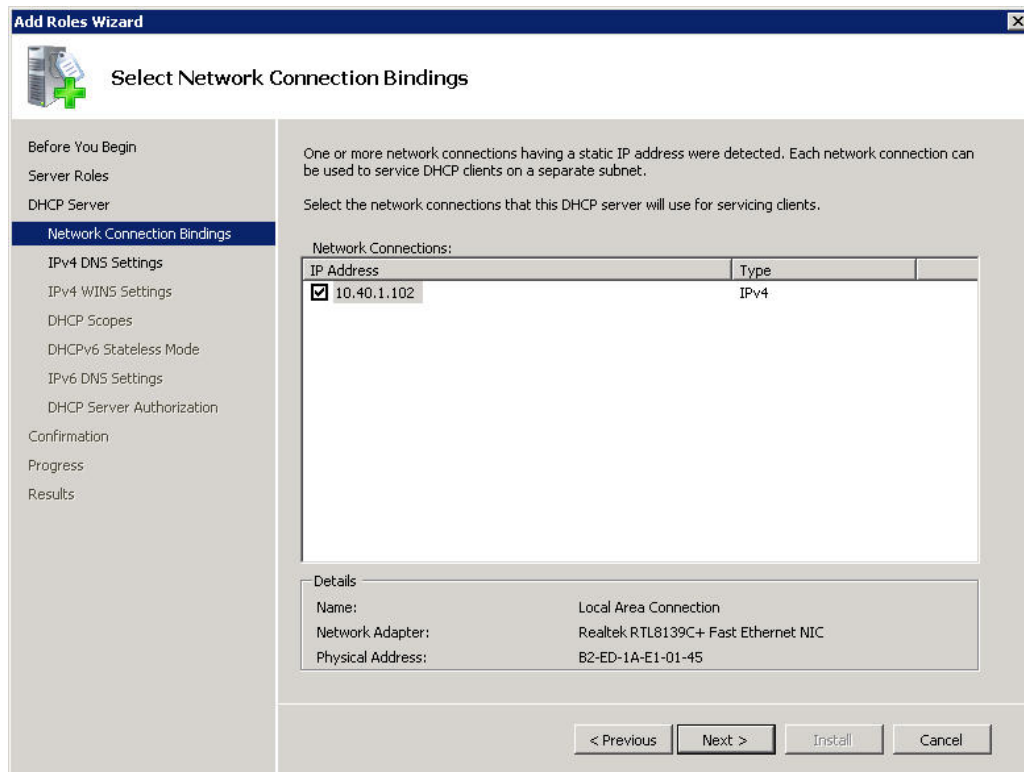
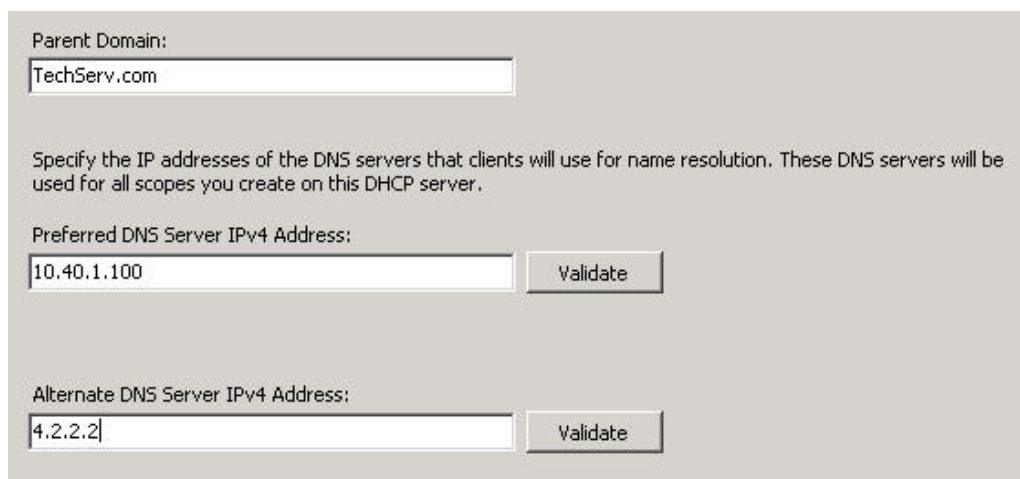


Figure 3: Configuring a DHCP Server Role

6. Click **Next**.
7. In the next screen, shown below, make sure the parent domain is the name of your domain and specify the DNS address of your DNS server (usually your domain controller) and an alternate DNS server, if one is available. If not, you can use a free, public DNS server. Click **Next**.



Parent Domain:
TechServ.com

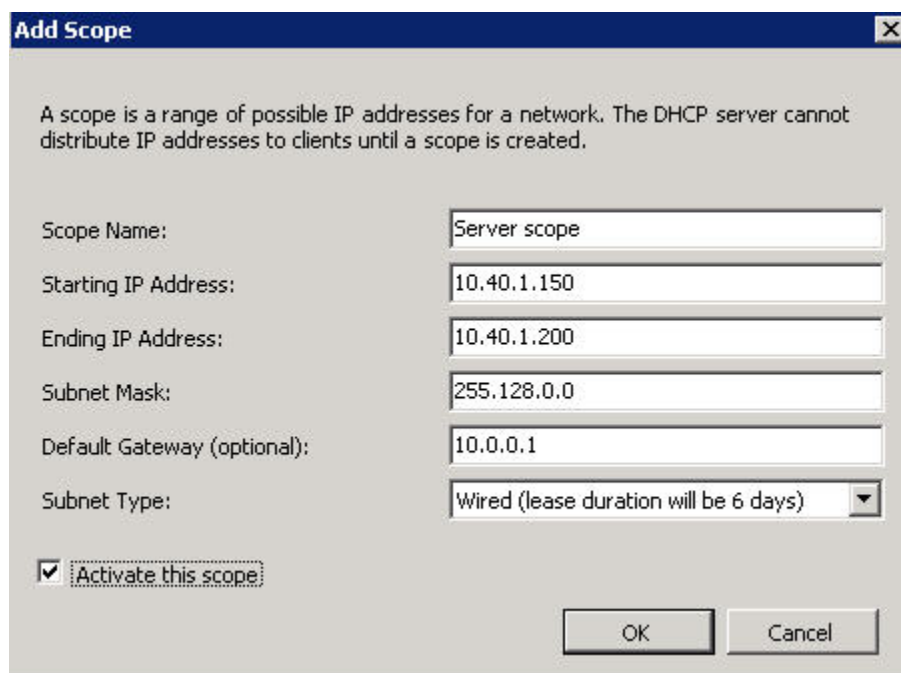
Specify the IP addresses of the DNS servers that clients will use for name resolution. These DNS servers will be used for all scopes you create on this DHCP server.

Preferred DNS Server IPv4 Address:
10.40.1.100

Alternate DNS Server IPv4 Address:
4.2.2.2

Figure 4: Entering DNS Information

- In the next box, select whether or not you would like to use **WINS** and click **Next**. Unless you know you'll need WINS for your network, we strongly recommend clicking "WINS is not required for applications on this network."
- Under the scopes section click **Add**. Enter the range of IP addresses that you'd like the DHCP server to offer. Click **Next**.



Add Scope

A scope is a range of possible IP addresses for a network. The DHCP server cannot distribute IP addresses to clients until a scope is created.

Scope Name: Server scope

Starting IP Address: 10.40.1.150

Ending IP Address: 10.40.1.200

Subnet Mask: 255.128.0.0

Default Gateway (optional): 10.0.0.1

Subnet Type: Wired (lease duration will be 6 days) ▼

Activate this scope

Figure 5: Adding IP Address Scopes

- Click **OK**, then click **Next**.

11. In DHCP mode, leave **stateless mode** enabled, unless you would like to use DHCP to issue IPv6 addresses. Click **Next**.
12. On the following screen, keep the default settings unless you would like to use different credentials to authorize your DHCP server. (Note that this step can only be performed by a Domain Administrator.) Click **Next**.
13. Click **Install**.
14. Click **Finish** when the process completes.

On Server Core, Without a Graphical Interface

DHCP can be enabled through **Server Core** by executing the following command:

```
start /w ocsetup DHCPServerCore
```

Follow the onscreen instructions to configure a scope and then authorize it.

To uninstall the DHCP server role, enter the following command:

```
start /w ocsetup DHCPServerCore /uninstall
```

The DHCP server role can be started by entering the following commands:

```
sc config dhcpserver start= auto
net start dhcpserver
```

DHCP Options

A DHCP server can pass much more information than just an IP address to its clients. It can, for example, also inform the clients of the DNS server address, or about where to find the Network Time Protocol servers. Options can be configured at both the Scope and Server levels.

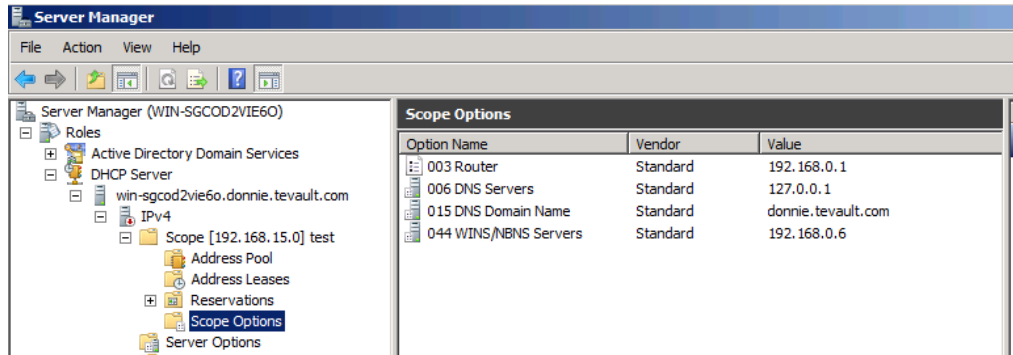


Figure 6: DHCP Scope Options

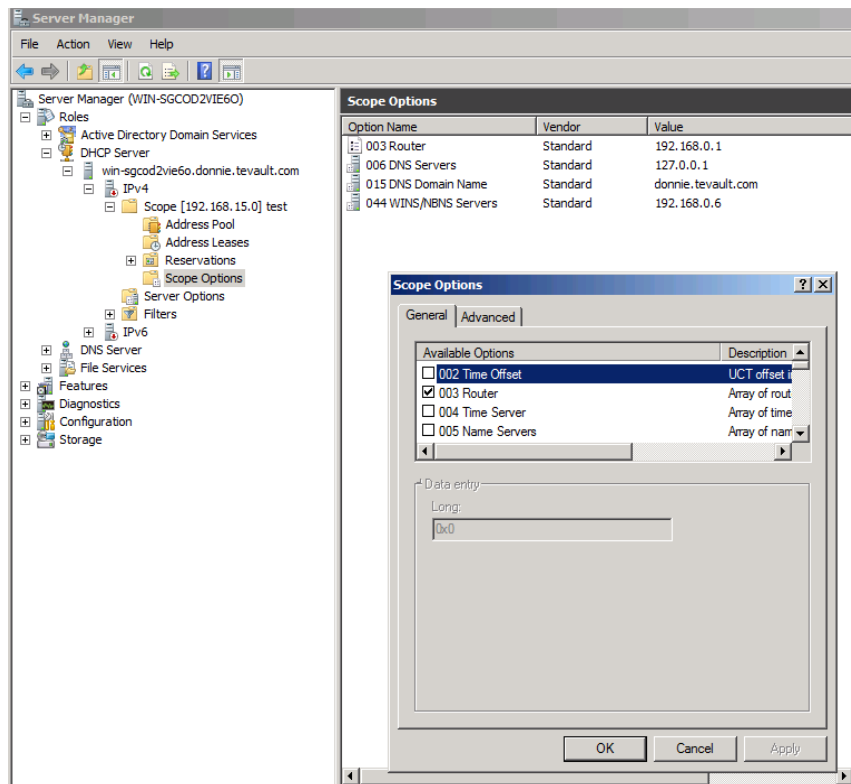


Figure 7: DHCP Scope Options

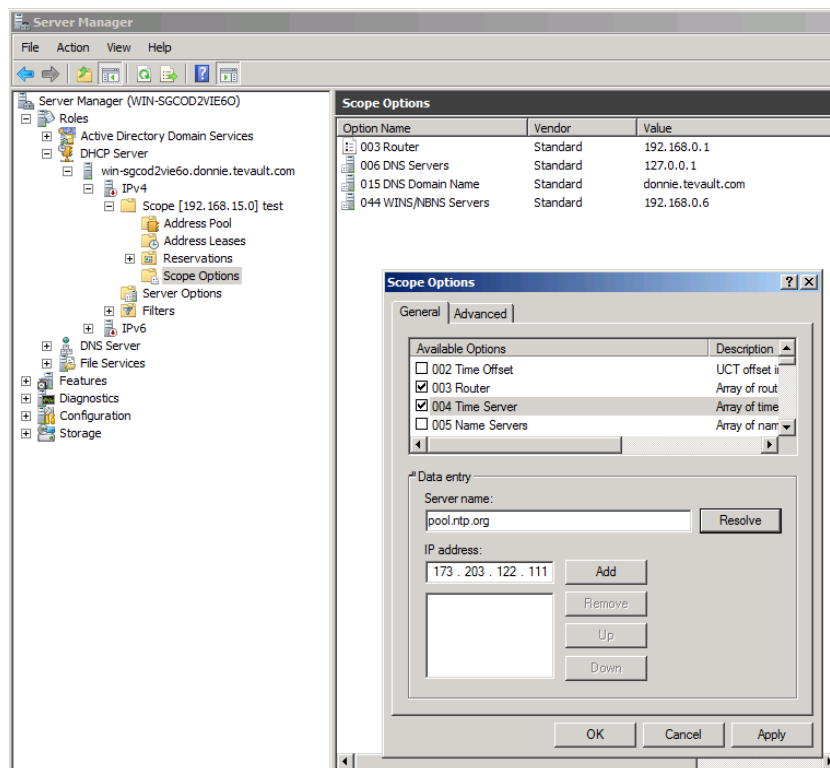


Figure 8: DHCP Scope Options

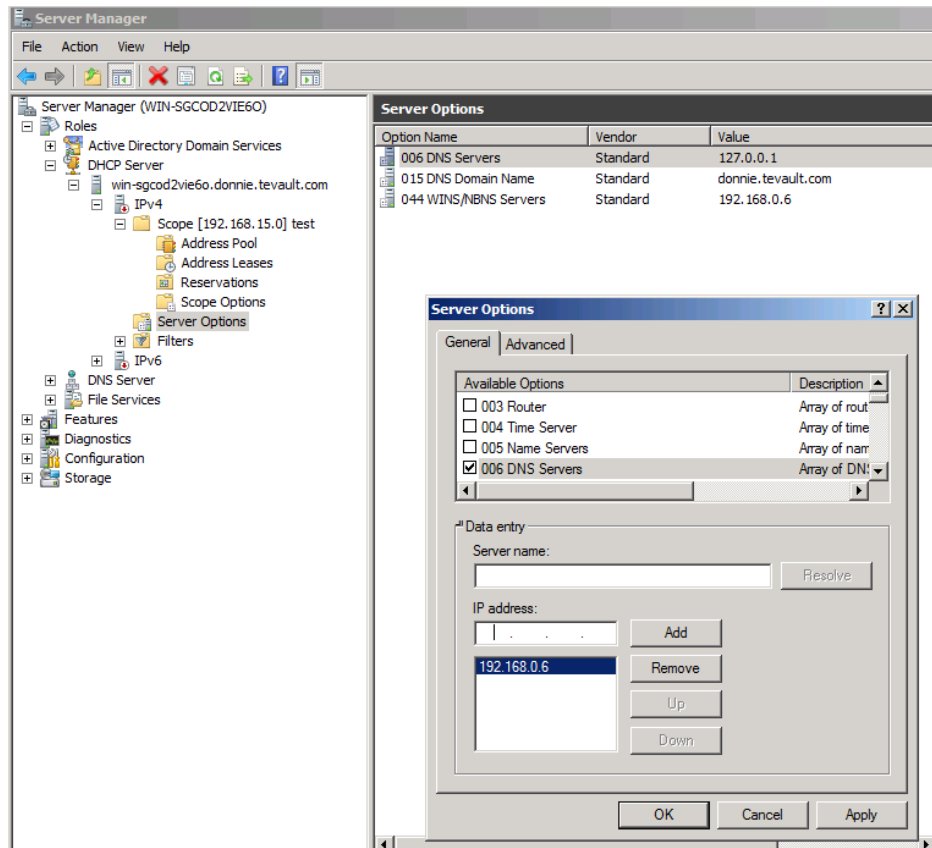


Figure 9: DHCP Server Options

DHCPv6

Normally, you would not configure DHCP to issue IPv6 addresses. Rather, DHCPv6 hosts would normally run in “stateless” mode. This means that a network host will issue itself an IPv6 address that’s compatible with the subnet on which it’s located. It calculates this address by exchanging Router Solicitation and Router Advertisement messages with its neighboring IPv6 router.

Certain options can be configured for IPv6.

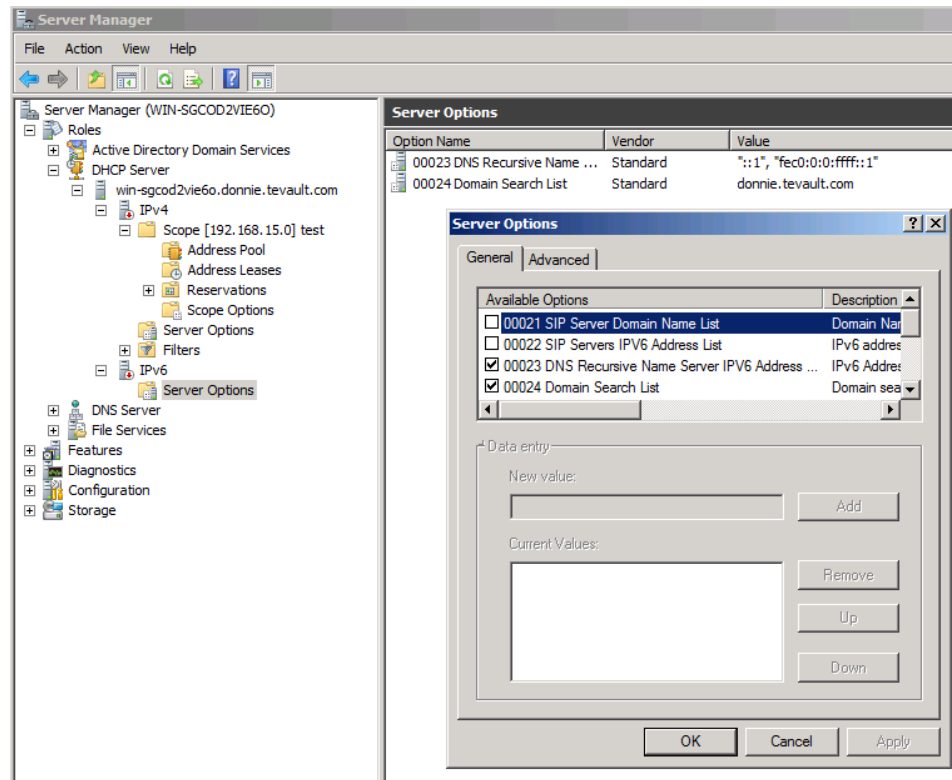


Figure 10: DHCPv6 Options

PXE Boot

The **Pre-Boot Execution Environment**, or PXE Boot, is a configurable Windows Server 2008 option that allows you to boot a host from a network service. To enable this, you must have a DHCP server running, as well as Windows Deployment Services. Typically, Microsoft will not require you to set this up on the 70-642 exam, because it involves Windows Deployment Services, which are mostly covered in the 70-643 exam. For this exam, you will just need to be familiar with the process, which is as follows:

1. A client sends a PXE request.
2. The PXE server forwards the request to a deployment services provider.
3. The PXE provider inspects the server and the address and either succeeds or fails.
4. If it fails, the process stops. If it succeeds, the request is fed to the deployment server and the boot process begins through the trivial file transfer protocol, or TFTP. TFTP is a simple file protocol that passes through port 20.

Note that if you have DHCP and Windows Deployment Services running together on the same machine, both services will try to listen on port 67. To rectify this, you'll need to modify a registry key to force WDS to listen on another port. You'll want to modify this key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WDS\Server\Providers\WDSPXE

Set the key value to:

UseDHCPPorts = 0

Then open a command-prompt window and enter the command:

```
WDSUTIL /Set-Server /UseDHCPPorts:No /DHCPOption60:Yes
```

Relay Agents

DHCP relay agents are used to forward DHCP requests across routers that aren't configured to forward broadcast messages. This allows a single DHCP server to issue IP addresses to clients on multiple subnets. You can enable a DHCP relay agent by performing the following actions:

1. Install **DHCP**, as we explained, above.
2. Under **IPv4/6** in **Routing and Remote Access**, right click **General** and click **New Routing Protocol**.
3. Select **DHCP Relay Agent** from the dropdown menu.
4. Right click the agent that is created underneath IPv4 and select **New Interface**.
5. Choose the interface you'd like to use as a relay agent, and then dynamically assign it an IP from the menu.

Scopes

A DHCP scope is a range of IP addresses that's available for use by a DHCP server. In other words, a DHCP scope is the range of addresses that your DHCP server is allowed to issue. There can be more than one scope, or just one scope. When your DHCP server runs out of usable addresses in one scope, it will attempt to use addresses from another scope. You can define scopes in the DHCP MMC.

Exclusions

You can configure DHCP to exclude certain IP addresses from being issued to clients. For example, if the DHCP scope were set for addresses 192.168.0.50 through 192.168.0.200, you could exclude the address of 192.168.0.100 if that address is statically assigned to a server.

To do this, you just need to follow this procedure:

1. Navigate to the **DHCP MMC** from **Administrative Tools** in the **Start Menu**.
2. Double-click your **DHCP server**.
3. Double-click **IPv4**.
4. Right-click **Address Pool** and add a **New Exclusions Range**.
5. Enter the starting and ending IP address range that you would like to exclude, and then click **Close**.

Authorizing a DHCP Server in Active Directory Domain Services

DHCP servers have to be authorized by Active Directory Domain Services before they're allowed to issue IP addresses. In other words, an unauthorized DHCP server cannot issue IP addresses to a host in the domain. This helps prevent security problems that could result from someone placing a rogue DHCP server on your network. You can either authorize a DHCP server when you first install DHCP, or you can do it by navigating to the DHCP MMC and adding an authorization.

Note that only Domain Administrators can authorize a DHCP server.

DHCP Reservations

If desired, you can reserve certain addresses from your scope for certain network hosts. You could, for example, reserve a particular address for a network printer. In essence, this allows your printer to now have a static address.

To do this, you'll need to know the MAC address of the host for which you want to reserve an address. When you create the reservation, you'll enter the Reservation Name, the desired static IP address, and the MAC address of the device that will use that IP address. Then choose whether the host supports DHCP, BOOTP, or both. Optionally, you can enter a description of the host.

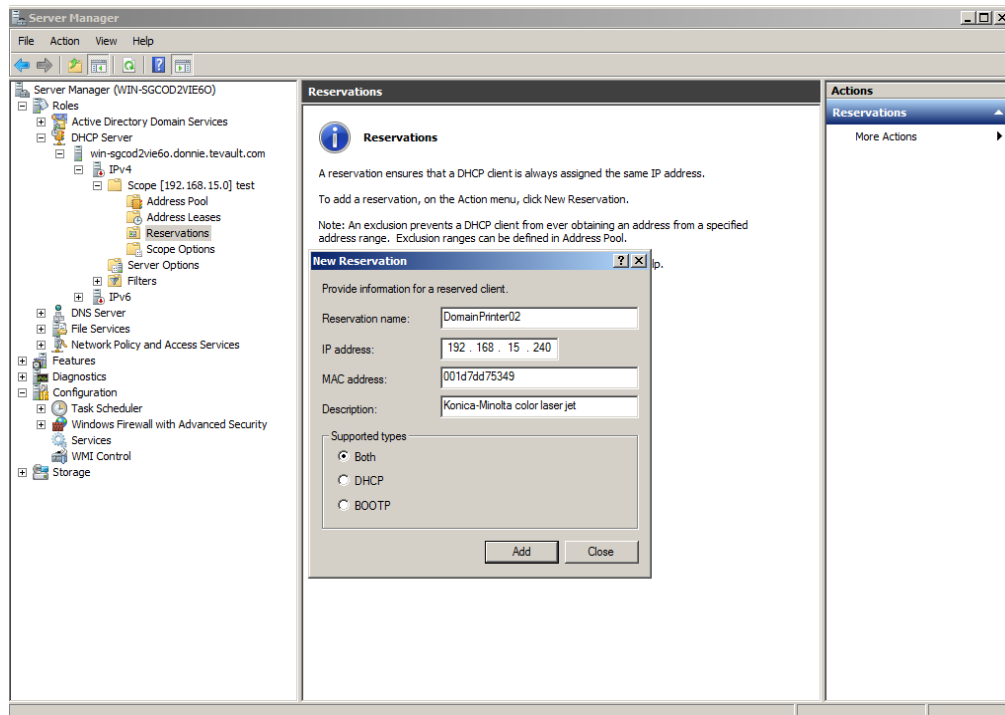


Figure 11

Note that if you have multiple DHCP servers serving a particular subnet, you'll want to have each reservation configured on each DHCP server. That way, if one DHCP server goes down, the reservation will still be available.

Windows Server Hyper-V

What you need to know about Windows Server 2008 Hyper-V and DHCP is pretty straightforward, although Microsoft will both have you think that it's really complicated, and also ask you a ton of questions on it. To get through it all, just know the following:

1. Hyper-V can run as a DHCP server.
2. Hyper-V can receive a DHCP address, as long as a network interface is assigned to the Virtual Machine.
3. Hyper-V supports IPv6.
4. Hyper-V plays friendly with Active Directory Domain Services domain building.

If they try to trick you on this, don't fall for it. It's actually completely straightforward.

Configuring Routing

When setting up a routing server, you can either set up static routes manually or use a routing protocol. For a small network, doing it the manual way might be fairly easy. On a larger network, it's advantageous to use a routing protocol that will discover the routes for you. Also, if a route were to change, for example, due to a failure of a network component along the route, using a routing protocol would allow the router to compensate automatically.

One way to use static routing is if you have a client machine with two different network interfaces for two different purposes. Let's say that one interface connects to the Internet, and the other interface connects to a remote subnet on the LAN. You would need to set up static routes in order to make both of these interfaces work properly.

Setting Up Static and Persistent Routes

To manually add a static route, open the command-prompt and enter something like:

```
route add 192.168.5.0 mask 255.255.255.0 192.168.5.1
```

The first IP address is the address of the destination network. The last IP address is the address of the gateway router that leads to that network. (Note that the gateway router must be on the same subnet as the host.)

Be aware that this route will disappear when the server is rebooted. If you want the route to be persistent after a reboot, just add the "-p" option switch to the command.

```
route add -p 192.168.5.0 mask 255.255.255.0 192.168.5.1
```

Using a Routing Protocol

To use a routing protocol, you'll need to install **Routing and Remote Access Services**. Next, you'll have to install Network Policy and Access Services, as shown in the steps below.

1. Start **Server Manager**.
2. Select **Roles**.
3. Click **Add Roles**.
4. Select **Network Policy and Access Services**.
5. Click **Next** at the intro screen.
6. Select **Routing and Remote Access Services**, and then click **Next**.
7. Click **Install**.
8. Once it's installed, click **Close**.

When deciding on which routing protocol to use, you need to understand the difference between classful and classless routing protocols. RIPv1, often simply referred to as "RIP," is a classful routing protocol that uses a hop-based metric system to count distances between routing points. RIPv1 works great for small networks, but doesn't scale well in large networks.

RIPv2 is a classless protocol, and is a much better choice for large networks.

To set up routing protocols (RIP):

1. Navigate to **Routing and Remote Access** from **Administrative Tools** and right-click on your server.
2. Select **Configure and Enable Routing and Remote Access**.
3. The Wizard will open; click **Next**.
4. Select **Custom Configuration**; click **Next**.
5. Select the **LAN Routing** checkbox; click **Next**.
6. Click **Finish**. If you're asked to start the service, start the service.
7. Your server should now have a green up arrow next to it. **Expand** your Server and **expand IPv4**, then right-click **General** and select **New Routing Protocol**. You can select either **RIPv2** or **IGMP**. Select **RIPv2**. **Note:** This is the same spot where you'd enable a DHCP relay agent.
8. Now you'll see RIP listed under IPv4. Right-click it, and select **New Interface...**
9. Select your interface.
10. Press **OK**.

NOTE: The OSPF routing protocol is no longer supported by Windows Server 2008.

RIPv2 vs. IGMP

Although you don't need to know much about routing protocols to pass the 70-642 exam, you may want to be aware of the differences between the two routing protocols that are supported by Windows Server 2008.

RIPv2, short for "Routing Information Protocol version 2", is a fairly simple protocol to set up. Its main disadvantage is that it's only useful for small networks. It uses "hop counts" as a routing metric, and to prevent "routing loops", it places a limit of 15 on the number of hops it can make in order to reach a remote host.

IGMP, short for "Internet Group Management Protocol", is used to establish multicast group memberships. When used for gaming and streaming video, IGMP makes for a more efficient use of resources.

Configuring IPSec

IPSec is one of the best available methods to secure network traffic. It provides many security benefits, including data encryption, data integrity verification, data authenticity verification, and anti-replay protection. The IPSec implementation that's built into Windows Server 2008 R2 is compatible with Network Address Translation-Traversal (NAT-T). Also, Windows Server 2008 R2 IPSec supports Diffie-Hellman encryption keys up to 2048 bits in length. This type of encryption key is considered to be virtually unbreakable.

In addition, you can use IPSec to create filter rules that will perform specific actions on different types of network traffic. For example, you can set up a rules list that will deny Telnet traffic from anywhere but the local subnet, and that will also encrypt Telnet data exchanges.

On Windows Server 2008, you can configure IPSec through either Group Policy or through Connect Security Rules. However, it's highly recommended that you always use Connect Security Rules when dealing with either Windows Vista or Windows 7 clients. For use with clients running Windows XP or older, you'll need to use Group Policy.

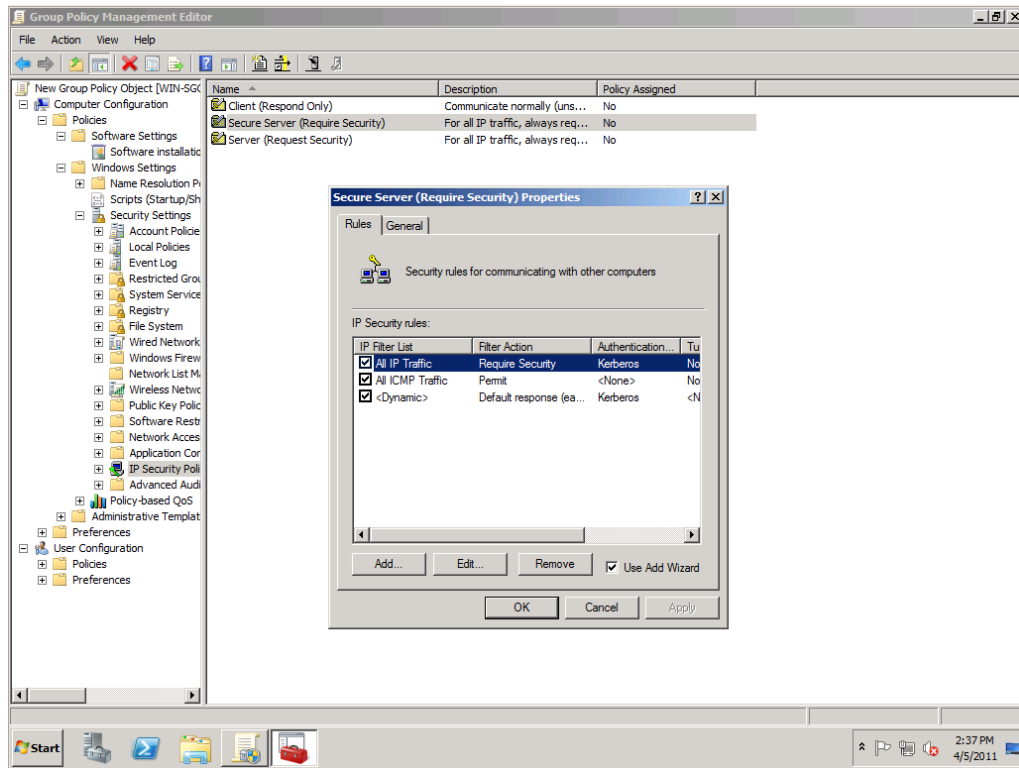


Figure 12: Setting IPSec Group Policy

To configure IPSec:

1. Navigate to the **Group Policy Management Console** (GPMC) by going to **Administrative Tools** and then **Group Policy Management**.
2. Select a group policy object and right-click it, then select **Edit**.
3. This will bring up the **Group Policy Management Editor**.
4. Expand **Windows Settings** and then **IP Security Policies on Active Directory**.
5. Create a new **IPSec Policy** by right-clicking and selecting **Create IP Security Policy**. This will bring up the IPSec Policy Wizard.

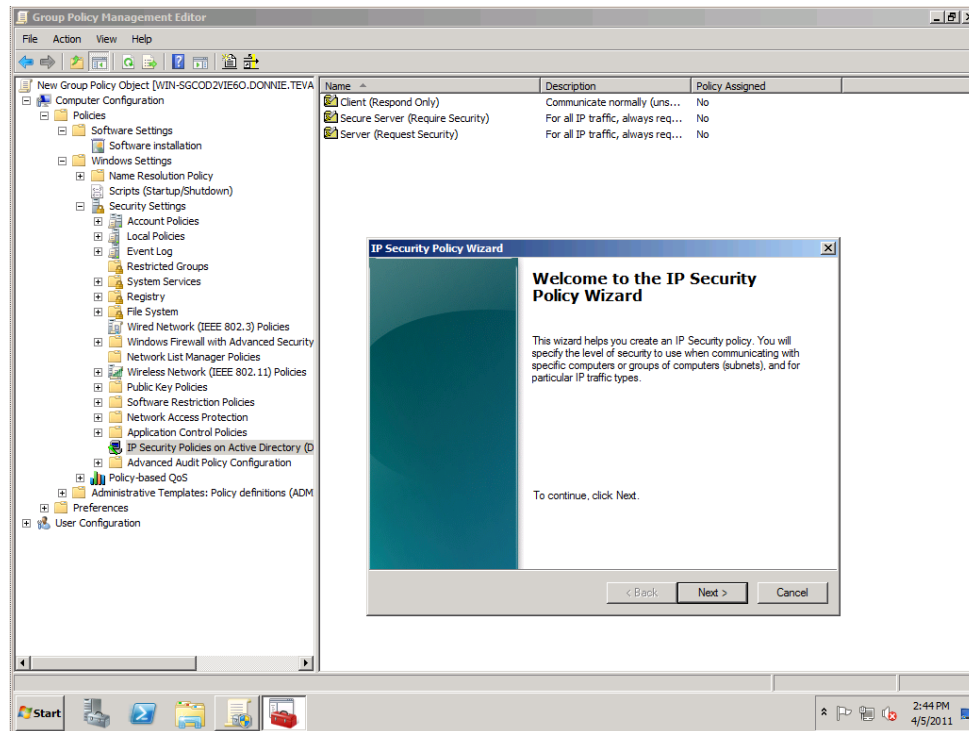


Figure 13: Configuring IPsec

6. Name the policy something appropriate.
7. Leave the default response for secure communication blank, unless you're dealing with external client authentication. Click **Next**.
8. From **IPsec Properties**, you can select **Edit** and choose from one of three options:
 - a. **Kerberos** – This is the default option for use within an Active Directory forest. When used within Active Directory, this is the easiest method.
 - b. **Certificates** – You can use certificates and their associated public/private encryption keys. Certificates can be issued by either an internal Certificate Authority, or by a commercial Certificate Authority.
 - c. **Preshared Keys** – Passwords are stored in plain-text files on each computer that's involved in an IPsec scenario. This option doesn't provide the same level of authentication security that Kerberos and Certificates do, and is recommended for use only on an internal test network.
9. To add more filters, click the **Add** button, and then click **Next** when the Wizard pops up.
10. You can either specify the rule for a specific tunnel in the IPsec policy or, if you do not wish to use an IP tunnel, select "This rule does not specify a tunnel."
11. Click **Next**.
12. Select **All Network Connections**, **LAN**, or **Remote Access**, depending on what type of filter you'd like to set up, and click **Next**.

13. On the next screen, shown below, you can filter traffic on either **ICMP** or **All IP Traffic**. Select the appropriate radio button and click **Edit**.

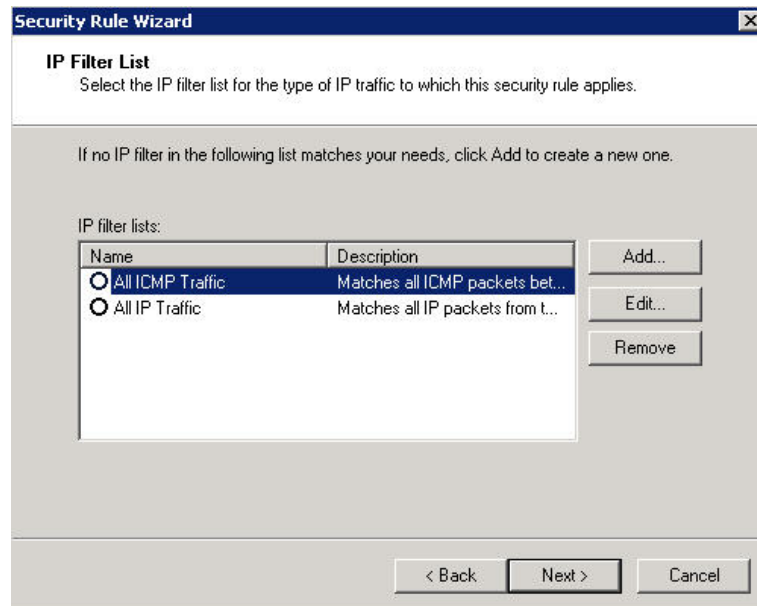


Figure 14: Creating an IPSec Filter

14. Name the filter something, such as **"RDP Traffic."**
15. Click **Edit**.
16. On the addresses page, you can specify addresses you'd like to restrict. The choices are My IP address, a specified address, DNS, WINS, DHCP, or Default Gateway. Choose appropriately.
17. On the protocol page, you can specify the type of traffic you'd like to filter. For example, you could choose RDP.
18. Click **OK**.
19. On the filter page, you can choose to **Permit**, **Request Security**, or **Require Security**. For example, if you wanted to restrict RDP, you could choose "Require Security."
20. Click **Next**.
21. At this point, you can choose Kerberos, Certificates, or Preshared Keys. Kerberos is the default choice, and is the most secure.

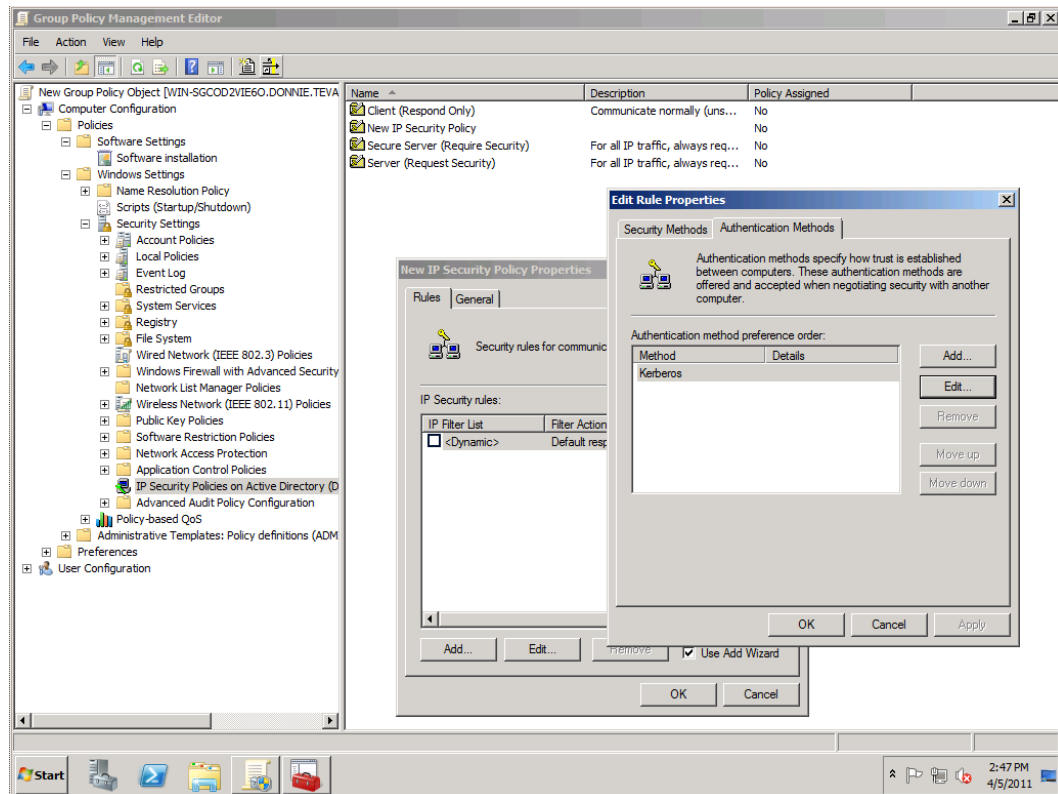


Figure 15: Setting IPsec Group Policy

22. Click **Next**, and then click **Finish**.

Security Associations

Whenever two computers negotiate an IPsec connection, they form a “Security Association,” or “SA.” The two protocols that are used to create Security Associations are the Authentication Headers (AH) protocol, and the Encapsulating Security Payload (ESP) protocol. These two protocols can be used separately, or both can be used together.

Authentication Headers

If the Authentication Headers protocol is used by itself, data origin authentication, data integrity verification, and anti-replay protection will be effected for the entire IP packet. However, the data payload is not encrypted.

Encapsulating Security Payload

The Encapsulating Security Payload protocol only works with the data payload. In addition to data origin authentication, data integrity verification, and anti-replay protection, ESP also encrypts the data payload.

NOTE: For the exam, be sure to remember this simple rule: If you need data encryption, use ESP. If you don't need data encryption, use AH.

Domain 2: Configuring Name Resolution

Name resolution is a vital part of any successful organization. Without proper name resolution, we would forever be required to remember IP addresses, instead of human-friendly names. Instead of typing "google.com" in a connection box, we would instead have to type out the full 32-bit address.

Name resolution is primarily handled via the Domain Name System, or DNS. A Domain Name System can run on Windows Server 2008, and will convert numerical IP addresses to dedicated, human-friendly names. Through DNS, the whole internet is connected through a fully qualified domain name system.

Take, for example, the address `www.preplogic.com`. This address consists of three parts:

- **Top Level Domain** – .COM
- **Domain** – preplogic
- **Sub-domain** – www

A host address that contains all three of these parts is called a "Fully Qualified Domain Name."

From a DNS perspective, this fully qualified Domain Name tells us:

- This server belongs to the commercial Domain Realm (.COM).
- It resides under the **preplogic** organization.
- The sub-domain should be listed as a `www` (World Wide Web) "A" record in DNS.

In this section, we're going to learn as much as possible about Windows Server 2008 and the Domain Name System. We will begin with configuring DNS.

Configuring a Domain Name System (DNS) Server

The Domain Name System server is configured through the Server Manager in the same way you would configure any other Server Role:

1. Open the Server Manager.
2. Select the DNS Server Role.
3. Click Next/Install.

As we'll see in a while, there are several different types of DNS servers that you can set up, with different types of DNS zones.

During the installation, pick a zone type in the menu screen.

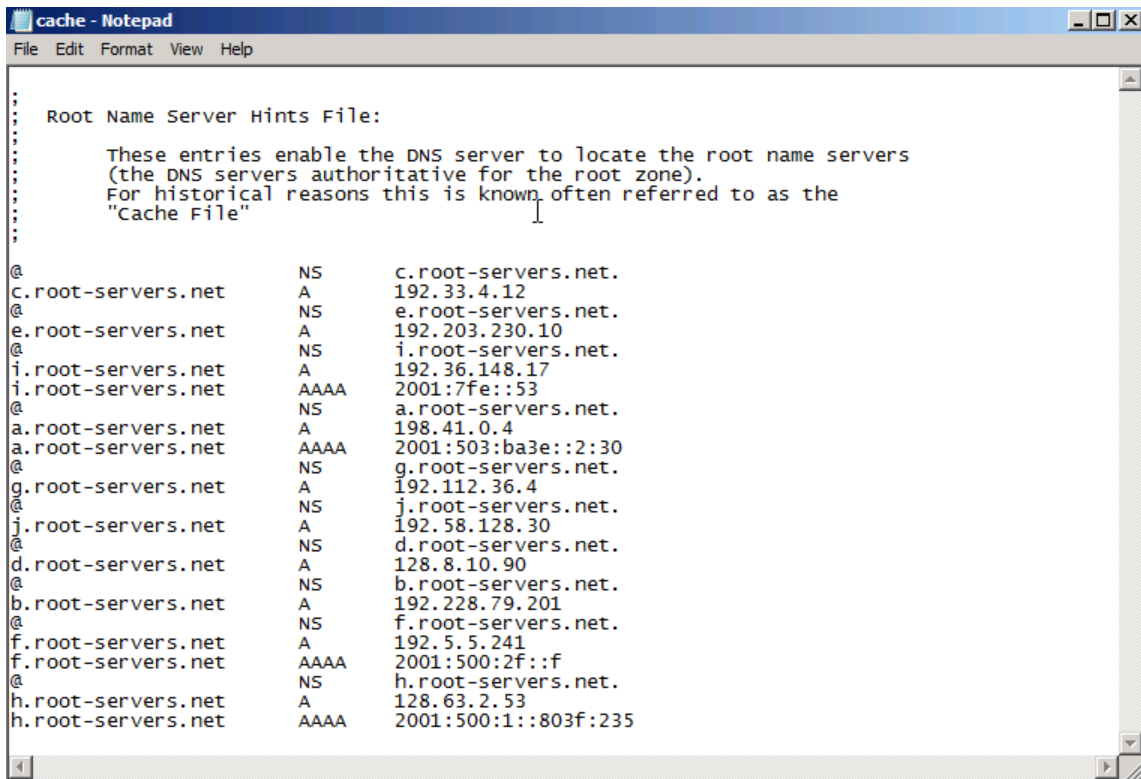
There, you will need to choose what type of DNS zone you will require. In terms of "installing" DNS, that's as complicated as it gets. However, there are other DNS specifics you will need to know about in order to successfully support it.

Root Servers and Root Hints

Your organization's DNS server will have information about hosts on the Local Area Network. However, if someone on the Local Area Network needs to access a site on the Internet, the local DNS server may not have that information stored locally. In that case, the local DNS server would either need to access a forwarding server, or know how to access the Internet root DNS servers. Information about the root DNS servers is stored in a "root hints" file.

A pre-configured root hints file, with information about the Internet root servers, is installed by default on Windows Server 2008 DNS servers. It's called the "cache.dns" file, and you'll find it in the "Windows\System32\Dns" directory. This file is a plain-text file that can be edited with any text editor. You can also edit it with the Windows Server DNS snap-in.

On a private network where Internet access isn't required, you can have root DNS servers that only contain resource records for hosts on the private network. In a case such as this, you can replace the root hints files on the lower-level DNS servers with files that only point to the internal root DNS server. DNS servers that serve as root servers won't have a root hints file.



```

Root Name Server Hints File:

These entries enable the DNS server to locate the root name servers
(the DNS servers authoritative for the root zone).
For historical reasons this is known, often referred to as the
"Cache File"

@
C.root-servers.net      NS      c.root-servers.net.
                        A       192.33.4.12
@
e.root-servers.net      NS      e.root-servers.net.
                        A       192.203.230.10
@
i.root-servers.net      NS      i.root-servers.net.
                        A       192.36.148.17
i.root-servers.net      AAAA   2001:7fe::53
@
a.root-servers.net      NS      a.root-servers.net.
                        A       198.41.0.4
a.root-servers.net      AAAA   2001:503:ba3e::2:30
@
g.root-servers.net      NS      g.root-servers.net.
                        A       192.112.36.4
@
j.root-servers.net      NS      j.root-servers.net.
                        A       192.58.128.30
@
d.root-servers.net      NS      d.root-servers.net.
                        A       128.8.10.90
@
b.root-servers.net      NS      b.root-servers.net.
                        A       192.228.79.201
@
f.root-servers.net      NS      f.root-servers.net.
                        A       192.5.5.241
f.root-servers.net      AAAA   2001:500:2f::f
@
h.root-servers.net      NS      h.root-servers.net.
                        A       128.63.2.53
h.root-servers.net      AAAA   2001:500:1::803f:235

```

Figure 16: The "cache.dns" file

To look at the root hints with the DNS snap-in, do the following:

1. Navigate to **Start > Administrative Tools > DNS**.
2. Select your server.
3. Double-click **Root Hints**.
4. The following screen will appear:

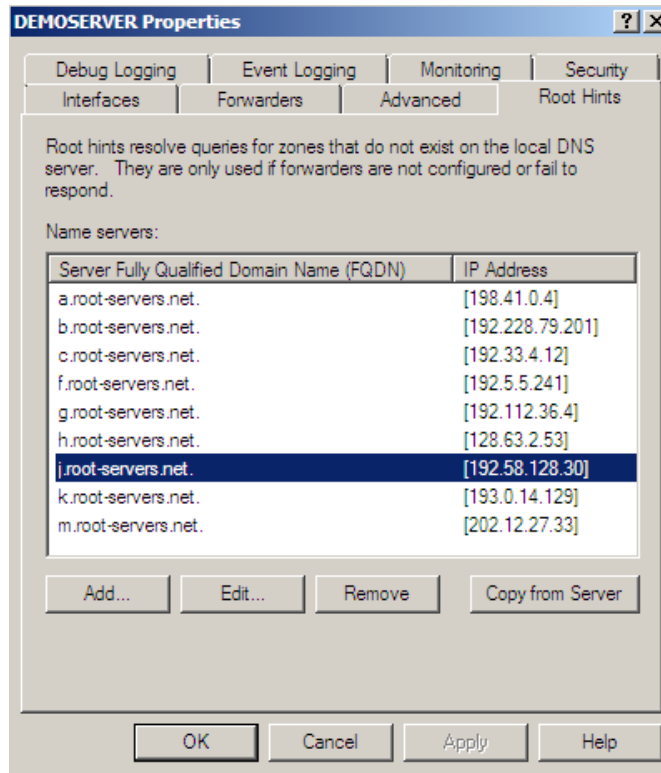


Figure 17: Adding a Root Hint, Part 1

5. To add a Root Hint, select **Add...**
6. In the screen below, enter a **Fully Qualified Domain Name** and the **IP Address**.

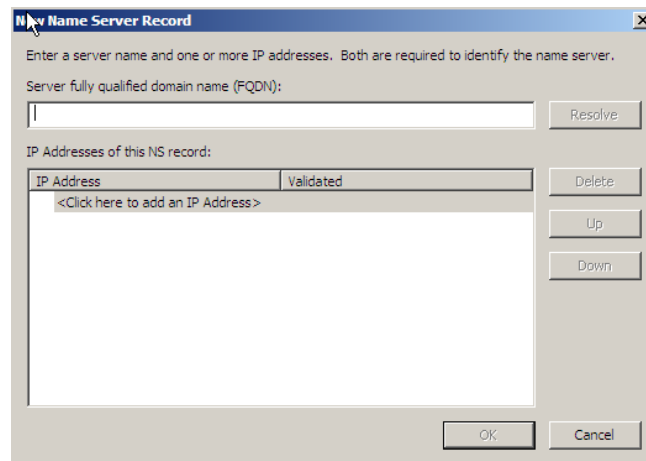


Figure 18: Adding a Root Hint, Part 2

7. Click **Resolve**.
8. Click **OK**. The name will show up in your root hint list.

For a DNS server that's used to access the Internet, you should never have to manually add or delete root hints. On a private LAN that doesn't require Internet access, you can delete the root hints for the Internet root servers, and add the appropriate root hints for the top-level DNS servers in your organization's domain.

Note also that setting up a DNS server to use both a forwarding server and root hints can give you a bit of fault tolerance. If the forwarding server goes down, you can have your DNS server automatically switch over to using its own roots hints.

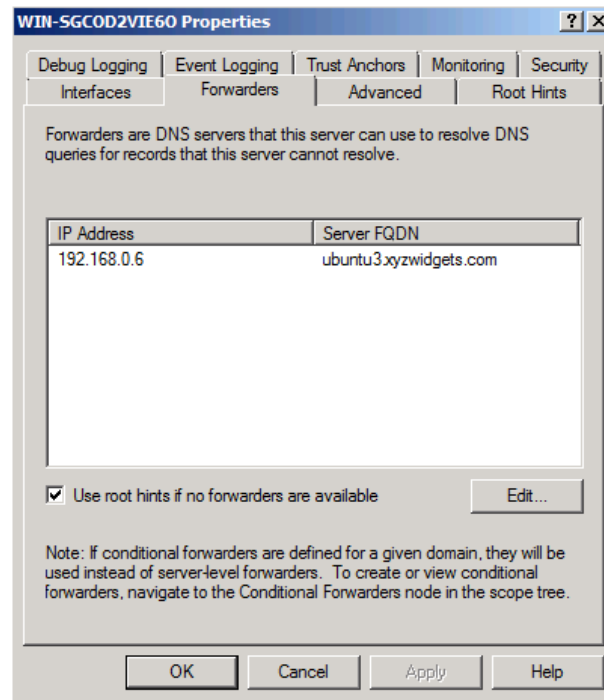


Figure 19: Using a Forwarder with Root Hints

Caching-Only

A caching-only DNS server has no zone information of its own. Every name resolution request that it receives is passed to another DNS server. The caching-only server then retains this look-up information on its local hard drive for a given period of time. This can help make DNS look-ups much faster, and can help conserve bandwidth at branch offices. Setting up a caching-only DNS server requires only that you not configure any zones on the DNS server, and that you verify that a root hints file exists on that server.

Conditional Forwarding

Conditional Forwarding with DNS is a basic if-then statement for name resolution. This permits name resolution queries for specific domains to be forwarded to certain specific DNS servers. An example of how this would be handy is if two companies were to merge with each other. Each company could then set up a server to forward DNS queries to a server on the other company's network.

To add a conditional forwarder, do the following:

1. Navigate to **Start Administrative Tools > DNS**.
2. Select **Conditional Forwarders**.
3. Right-click the **white space** in the Conditional Forwarders window and select **New Conditional Forwarder**.
4. On the screen you see below, specify the domain you want the conditional forwarder to apply to, and then select the IP address of the DNS server to use in order to resolve that name.

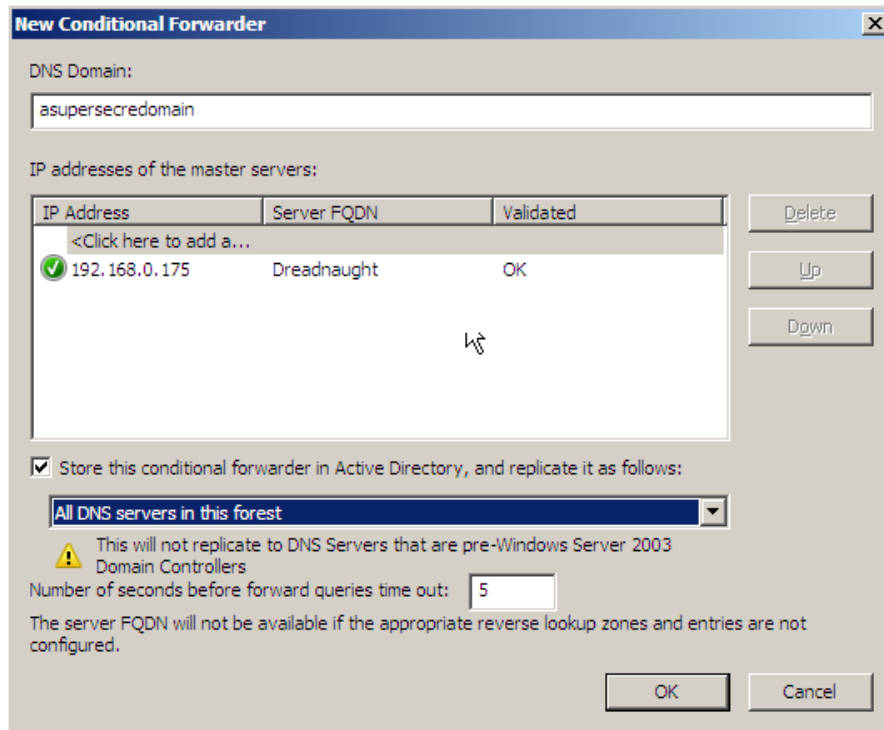


Figure 20: Adding a Conditional Forwarder

5. If you wish to store the conditional forwarder in Active Directory Domain Services, click the check box: "Store this conditional forwarder in Active Directory, and replicate it as follows:"
6. From the drop-down, choose either **All DNS servers in the forest**, **All DNS servers in this domain**, or **All domain controllers in this domain** (for Windows 2000 compatibility).
7. Click **OK**. The conditional forwarder will appear in the white space.

DNS Round Robin

By selecting the "DNS Round Robin" option, you can spread the workload out across multiple servers. Let's say that you have three web servers, all participating in a round robin and all serving out of the same website. When a request for that site comes from one client, the DNS server will send it to Web Server 1. When a request comes from a second client, it will be sent to Web Server 2. The third client request will get sent to Web Server 3. The cycle will start over with request number 4. This can increase website performance by helping to ensure that no web server gets bogged down with too much traffic.

Netmask Ordering

The netmask ordering option is another handy feature when there are multiple servers available. Let's again say that you have three web servers on your Wide Area Network (WAN), all serving the same website. However, each of the three servers is on a different segment of the WAN. With netmask ordering enabled, the DNS server will look at the IP address of the destination server and compare it with the IP address of the requesting client. The server request will then be directed to the web server that is nearest to the client. This is especially helpful if a server is on the same WAN segment as the client, as this will help reduce WAN bandwidth usage.

WINS and Windows Server 2008

If you choose to enable NetBIOS name resolution, you might consider setting up a WINS server. WINS isn't a necessity if you're only dealing with one subnet, but it's a must if you want to enable NetBIOS across multiple subnets.

To set up WINS, you will need to integrate it as a feature of your DNS server. Do so by completing these steps:

1. Open **Server Manager**.
2. Click **Features**.
3. Choose **Add A Feature**.
4. Select **WINS Server**.
5. Click **Install**.
6. Click **Close** once the install completes.

Based on the configuration you require, you will need to choose what is appropriate for your configuration.

Configuring DNS Zones

Hosts on a network can be categorized into DNS zones. Each individual zone contains authoritative information about a specified portion of the DNS namespace. A zone can have a name such as "preplogic.com." In turn, this zone can have subzones, such as "tampa.preplogic.com" and "stmarys.preplogic.com."

There are several different types of zones that you can set up. You'll want to be familiar with each one.

Types of Zones

Primary Zones

On DNS servers that are *not* integrated into Active Directory Domain Services, Primary Zones contain the original read-write source data for a particular portion of the DNS namespace. They can contain subzones and usually hold host record types. When an address is updated or changed, it is updated on a server that hosts a primary zone. Primary zone data are stored in plain-text files, and are not stored in Active Directory Domain Services.

Secondary Zones

Secondary Zones are replicated from another server. They don't contain their own local data, but instead are read-only copies of data from other servers. When the primary zone is updated, the servers that host secondary zones will receive the updates via a zone transfer.

The main reason to use secondary zones is to provide load relief for higher bandwidth DNS servers and to provide fault tolerance. If a server that hosts a primary DNS zone goes down, a server that hosts a secondary zone can take over.

If a record in the primary zone gets changed and it's important to transfer it to the secondary zone immediately, you can accomplish this by running the following command on the secondary server:

```
dnscmd /zonerefresh
```

Secondary zone data are stored in plain-text files, and are *not* stored in Active Directory Domain Services.

NOTE: The concept of Primary Zones and Secondary Zones only applies to DNS servers that have *not* been integrated into Active Directory Domain Services.

Stub Zones

Stub Zones only contain whatever resource records are necessary to find an authoritative DNS server for a master zone. They can be used to make DNS administration easier, and to improve name resolution performance. Also, stub zones can be used in a parent zone to maintain a list of current DNS servers in a child zone.

Active Directory Domain Services Integrated Zones

When you set up DNS on an Active Directory Domain Services domain controller, the default is for DNS to automatically be integrated into Active Directory. When you set up a DNS zone on a machine that isn't a domain controller, you have the choice of integrating it into Active Directory. With Active Directory Domain Services integration, the DNS zone is replicated to all DNS servers in either the domain or the forest, or to DNS servers that have an Active Directory partition. This way, if a DNS server exists, other DNS servers will be made aware of it and use it if necessary.

Additionally, you can create an Active Directory Domain Services partition for your DNS zone data that replicates all DNS data to the domain controllers and DNS servers throughout your network. To create a DNS Active Directory Domain Services partition, you can right-click your server and choose **Create Default Application Directory Partition**.

NOTE: The concept of Primary Zones and Secondary Zones does *not* apply to DNS servers that have been integrated into Active Directory Domain Services. That's because with Active Directory Domain Services integrated zones, all DNS servers are Primary servers.

GlobalNames Zones

Sometimes, having to use a whole Fully Qualified Domain Name to reference another computer on the network can be a bit awkward. One traditional way to get around this has been to continue using WINS and NetBIOS, even when DNS is available. However, WINS and NetBIOS are not compatible with IPv6, and their use will be fully deprecated when the transition to IPv6 is finally complete.

The GlobalNames zone is a feature that's brand-new in Windows Server 2008, and that works with IPv6. It allows DNS clients in an Active Directory forest to use a single-label name tag to connect to some specific server that's also located in the same Active Directory forest. You could, for example, simply use a name such as Fileserver to connect to a specific server. This is all done without the use of WINS or NetBIOS.

On a default installation of Windows Server 2008 DNS, the GlobalNames zone doesn't exist. In order to use this feature, you'll have to create the GlobalNames zone yourself.

You can then add servers to your GlobalNames zone by adding CNAME records to it. To create a GlobalNames zone, do the following:

1. Open **Command Prompt**.
2. Type **dnsmcd <your server name> /config /Enableglobalnamesupport 1**.
3. Open DNS by **Administrative Tools>DNS**.
4. Right-click your **DNS server**, choose **new zone**.
5. Click **Next**.
6. Choose **Primary Zone**.
7. Click **Next**.
8. Choose **Forward Lookup Zone**; click **Next**.
9. Name the zone **GlobalNames** and click **Next**.
10. Create a new file with this name: **GlobalNames.dns**; select **Next**.
11. Do not allow dynamic updates.
12. Click **Finish**.
13. Open **Command Prompt**.
14. Type **dnsmcd <your server name> /Zoneadd GlobalNames /Dsprimary /DP /forest**.

You can then add records to your GlobalNames zone by adding CNAME records to it. Also, note that the command in Step 2 above must be run on each DNS server that will need to resolve GlobalNames, even if the GlobalNames zones will be replicated to them.

Aging and Scavenging

Aging involves tracking the age of dynamically registered resource records. This is done by placing timestamps on those records. Scavenging involves deleting outdated resource records so that they don't accumulate over time. Aging and scavenging are both disabled by default. To use them, you'll need to enable them at both the server level and at the zone level.

There are two aging and scavenging settings with which you should be familiar. The No-Refresh interval is the time between the most recent refresh of a timestamp and the moment when the timestamp may be refreshed again. The Refresh interval is the time between the earliest moment when a timestamp can be refreshed and the earliest moment when the record can be scavenged.

Enabling scavenging is easy. All it requires is to click on a check-box.



Figure 21: Scavenging Stale Records

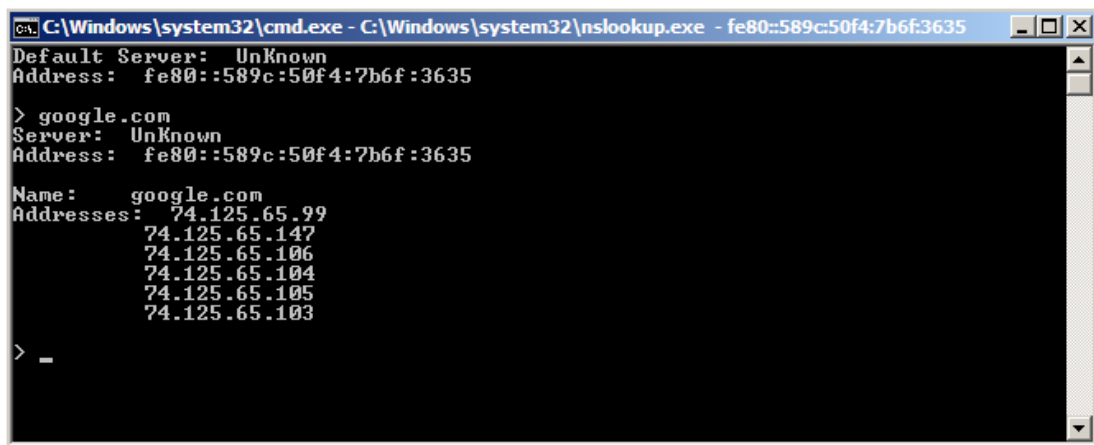
Clearing the DNS Server Cache

Let's say that the IP address for a certain website has changed, and that the change has already propagated to the DNS servers on the Internet. If clients from your organization try to visit that website, it could be that they won't be able to reach it. That could be due to having the old site information cached on the organization's DNS server. You can clear the cache on a DNS server by opening a command-prompt and entering:

```
dnscmd /clearcache
```

NSLOOKUP

Nslookup is a command-line tool that can be used to look up information regarding name resolution. From within DNS, you can launch nslookup by right-clicking your server and selecting Launch nslookup. You could also open a Windows command-prompt window and type in nslookup. When the program opens, type in the name of a domain to find information about the domain DNS servers.



```
C:\Windows\system32\cmd.exe - C:\Windows\system32\nslookup.exe - fe80::589c:50f4:7b6f:3635
Default Server:  UnKnown
Address:  fe80::589c:50f4:7b6f:3635

> google.com
Server:  UnKnown
Address:  fe80::589c:50f4:7b6f:3635

Name:    google.com
Addresses:  74.125.65.99
           74.125.65.147
           74.125.65.106
           74.125.65.104
           74.125.65.105
           74.125.65.103

> _
```

Figure 22: The Nslookup Tool

As you can see in the above screenshot, you can enter any domain name you choose, and the addresses will resolve. This command will show both authoritative and non-authoritative DNS servers for the domain that you specify.

Dynamic DNS Updating

Network hosts that are DNS clients can both register and dynamically update their resource record information on a DNS server. DNS clients that are configured with a static IP address will attempt to update both their host records and their pointer records. DNS clients that are also DHCP clients will only attempt to update their host records. If a network is running with a workgroup instead of a domain, then the DHCP server will update the pointer record on behalf of the DHCP client whenever an IP address lease is renewed.

The zone in which clients attempt to either register or update themselves must be configured to accept dynamic updates. You can configure a zone to allow non-secure updates, which can come from any computer. You can also configure a zone to allow only secure updates. This means that only computers that are members of the Active Directory domain can register themselves, and updates for a particular computer can only come from the computer that performed the original registration.

NOTE: For some reason that we don't understand, Microsoft has apparently decided to call this feature "Dynamic DNS." We're not sure why, since there's already another use for that term. When studying for the exam, keep in mind that the term as Microsoft currently uses it refers to the automatic DNS updating mechanism that we've described here.

Types of DNS Records

DNS supports many record types. A DNS “record” is just a recording of what type of server is located at that particular IP address. Windows Server 2008 supports the following common types of records:

- **Start of Authority Records**, or SOA records, contain information about properties for the zone.
- **Host Records** come in two varieties, both of which associate a computer name to an IP address. “A” host records are for IPv4 hosts, and “AAAA” host records are for IPv6 hosts.
- **Alias Records**, also known as CNAME records, allow using more than one name for a particular host. This comes in handy if the actual host name is long or difficult to remember.
- **Mail Exchange Records**, also known as MX records, are for mail servers.
- **Service Location Records**, also known as SRV records, specify the locations of specific SRV-aware services in a domain.
- **Name Server Records**, or NS records, specify which DNS server is authoritative for the zone.
- **Pointer Records**, or PTRs, are records that convert IP addresses to names.

DNS also supports other record types, including these:

Record	NOTE	RFC Article	Description
A6	38	RFC 2874	Experimental IPv6 record
AFSDB	18	RFC 1183	AFS servers – experimental
DNAME	39	RFC 2672	Delegation name for IPv6
DNSKEY	48	RFC 4034	DNSSEC.bis DNS public key
DS	43	RFC 4034	DNSSEC.bis DNS public key signer
HINFO	13	RFC 1035	Host Information
ISDN	20	RFC 1183	ISDN dedicated address
KEY	25	RFC 2535	Public Key Identity – encryption
LOC	29	RFC 1876	GPS data
NAPTR	35	RFC 3403	Naming Authority Pointer Record
NS	2	RFC 1035	Name Server; Authoritative Name Server for the domain
NSEC	47	RFC 4034	Next Secure record
NXT	30		DNSSEC Next Domain record type
RRSIG	46	RFC 4034	DNSSEC.bis. Signed RRset
RT	21	RFC 1183	Through-route binding
SIG	24	RFC 2931//2535	DNSSEC
SOA	6	RFC 1035	Start of Authority
SPF	99	RFC 4408	Sender Policy Framework
SRV	33	RFC 2872	Services
TXT	16	RFC 1035	Text information associated with a name
WKS	11	RFC 1035	Well Known Services
X25	19	RFC 1183	X.25 address

Figure 23: DNS Record Types

Configuring DNS Replication

DNS replication allows the process of sharing DNS information from one DNS server to another. Configuring DNS replication is a fairly easy process. All you have to do is right-click your DNS zone and choose the **General** tab, which will show up like what you see below. (Note that this procedure is for Active Directory Domain Services-integrated DNS.)

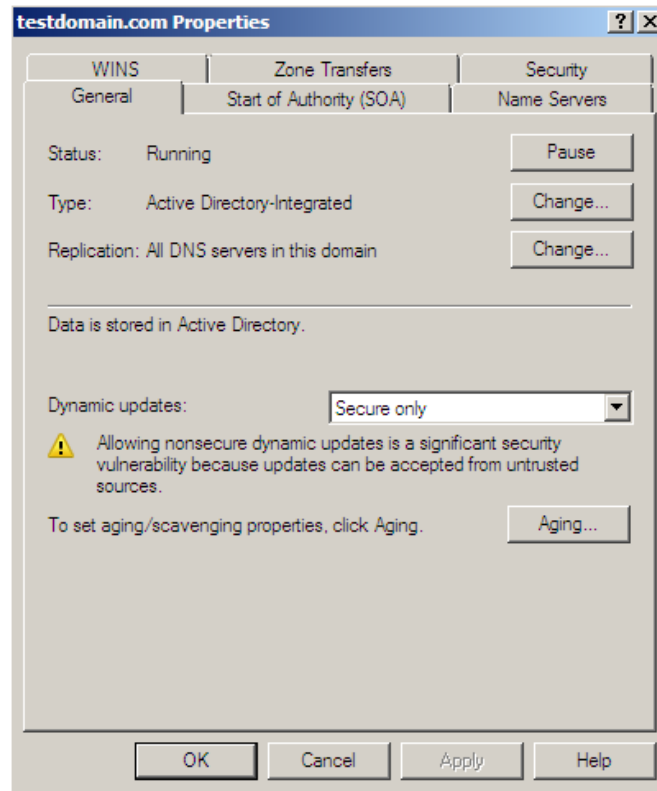


Figure 24: Setting Up DNS Replication

You can then click the change button. This will bring up the menu you see below.




Figure 25: Changing the DNS Replication Scope

DNSSEC

DNS is not an inherently secure system, and is susceptible to various forms of attack. Cache-poisoning attacks, spoofing attacks, and man-in-the-middle attacks are all possible and are all rather easily performed on DNS. DNSSEC can inform clients about whether DNS information is coming from a valid DNS server, about the integrity of the DNS data, and about whether a given host really does exist.

DNSSEC stores pairs of public and private encryption keys in certificates. These certificates are called "Zone Signing Keys," or ZSKs, and are used to sign DNS zones. In turn, these certificates are validated by another public and private key combination called the "Key Signing Key," or KSKs.

The public keys for the ZSKs and the KSKs are also stored in the DNSKEY record. This allows zone signatures to be validated.

The "Next Secure," or NSEC, record is used to prove the non-existence of DNS names. That way, if a record isn't retrieved from a DNS lookup, the requesting client can be positive that the record doesn't exist.

This signature of a DNS record is stored in the "Resource Record Signature," or RRSIG, record. Each "A" record and each NSEC record has a corresponding RRSIG record.

DS records, short for "Delegation Signer" records, confirm the validity of delegated DNS servers. This helps prevent man-in-the-middle attacks that result from placing a rogue DNS server where it could answer recursive lookups.

NOTE: The DNSSEC implementation that's built into Windows Server 2003 and the original version of Windows Server 2008 are based on a standard that is now obsolete. This older implementation is not compatible with the DNSSEC implementation that's built into Windows Server 2008 R2.

DNS Socket Pooling

"Socket Pooling" is what Microsoft calls its implementation of Source Port Randomization. This feature is new for Windows Server 2008 R2.

Having a recursive DNS server always contact an authoritative DNS server via a random output port, rather than by the same port every time, makes it more difficult for an attacker to redirect recursive lookups to a counterfeit authoritative server. In turn, this makes it harder to redirect clients to counterfeit websites.

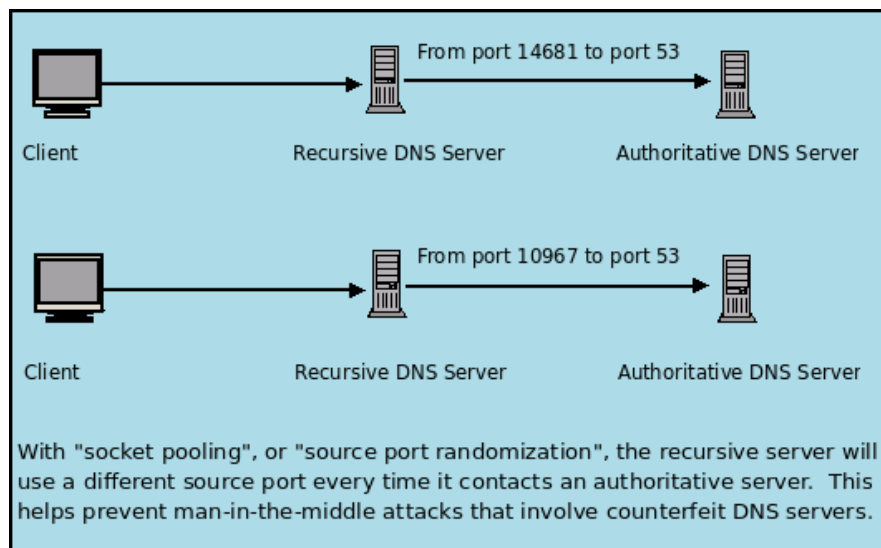


Figure 26: DNS Socket Pooling

DNS Cache-Locking

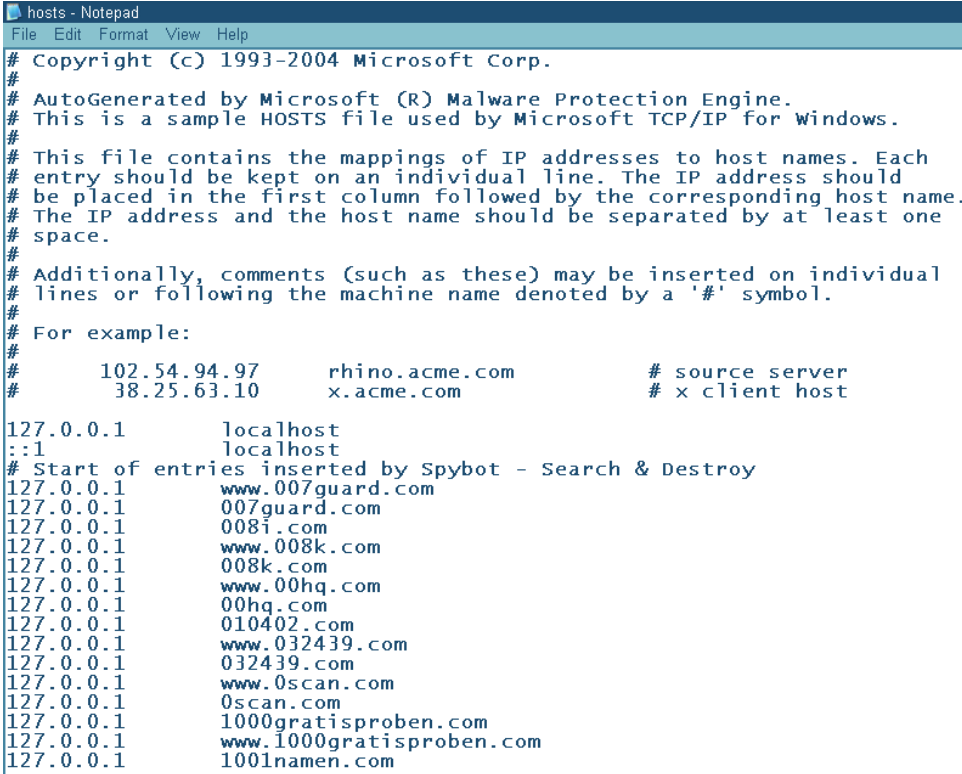
By enabling DNS cache-locking, you can prevent entries in a DNS server's cache from being overwritten until after the Time-to-Live value for that entry expires. This helps prevent attackers from deploying a cache-poisoning attack against your servers. This feature is also new in Windows Server 2008 R2.

Configuring Name Resolution for Clients

When a DNS server is available, each client on the network should be configured to use it. However, clients also have other ways to perform name resolution, which can be used either in place of or along with DNS.

The Hosts File

The hosts file can be used for name resolution when DNS isn't installed. Also, certain anti-malware programs can use it to prevent users from visiting known malware-contaminated websites. The hosts file is located at C:\Windows\System32\drivers\etc and it is simply called "hosts."



```
hosts - Notepad
File Edit Format View Help
# Copyright (c) 1993-2004 Microsoft Corp.
#
# AutoGenerated by Microsoft (R) Malware Protection Engine.
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com           # source server
#       38.25.63.10      x.acme.com              # x client host

127.0.0.1       localhost
::1            localhost
# Start of entries inserted by Spybot - Search & Destroy
127.0.0.1       www.007guard.com
127.0.0.1       007guard.com
127.0.0.1       008i.com
127.0.0.1       www.008k.com
127.0.0.1       008k.com
127.0.0.1       www.00hq.com
127.0.0.1       00hq.com
127.0.0.1       010402.com
127.0.0.1       www.032439.com
127.0.0.1       032439.com
127.0.0.1       www.0scan.com
127.0.0.1       0scan.com
127.0.0.1       1000gratisproben.com
127.0.0.1       www.1000gratisproben.com
127.0.0.1       1001namen.com
```

Figure 27: A Hosts File

Selecting a DNS Server

To select an appropriate DNS server, open the **Network Sharing Center**. Click TCP/IP, choose properties, and enter your DNS server information. (Note that this step isn't necessary for clients that have been configured to use DHCP.)

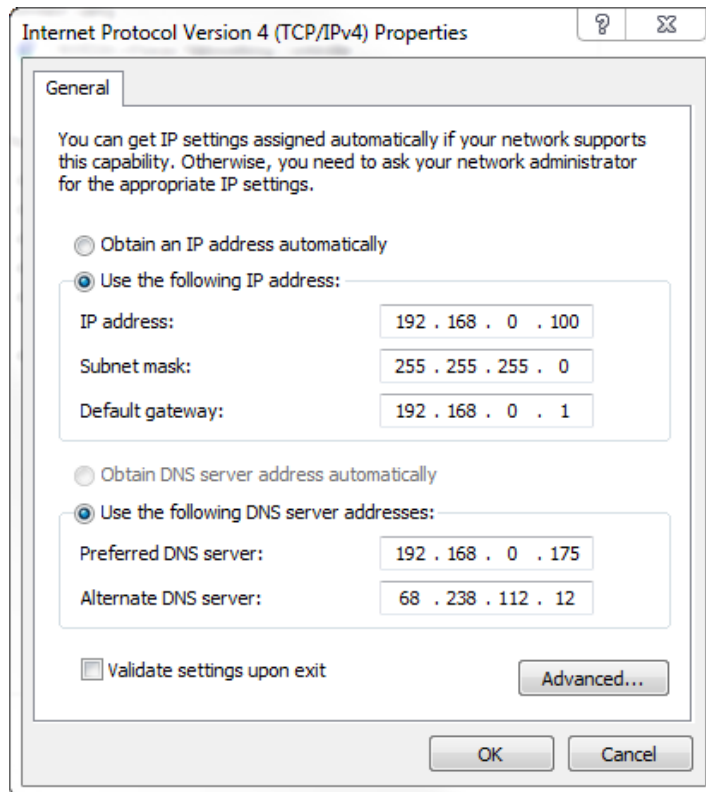


Figure 28: TCP/IPv4 Properties

LMHOSTS

Lmhosts is very similar to hosts, but is used with NetBIOS. The lmhosts file is located in the same place as the hosts file.

Managing Client Settings through Group Policy & Suffix Search Order

The clients on your network will either be statically configured to point to the correct DNS server, or they will receive DNS server information dynamically via the network DHCP server. You can manage other client DNS settings from the Group Policy Manager on your network's domain controller. For example, you can configure a group policy that will distribute a DNS suffix list to all DNS clients in the domain. This will allow users to access other computers by simply specifying a single-label computer name, rather than typing in the whole Fully Qualified Domain Name.

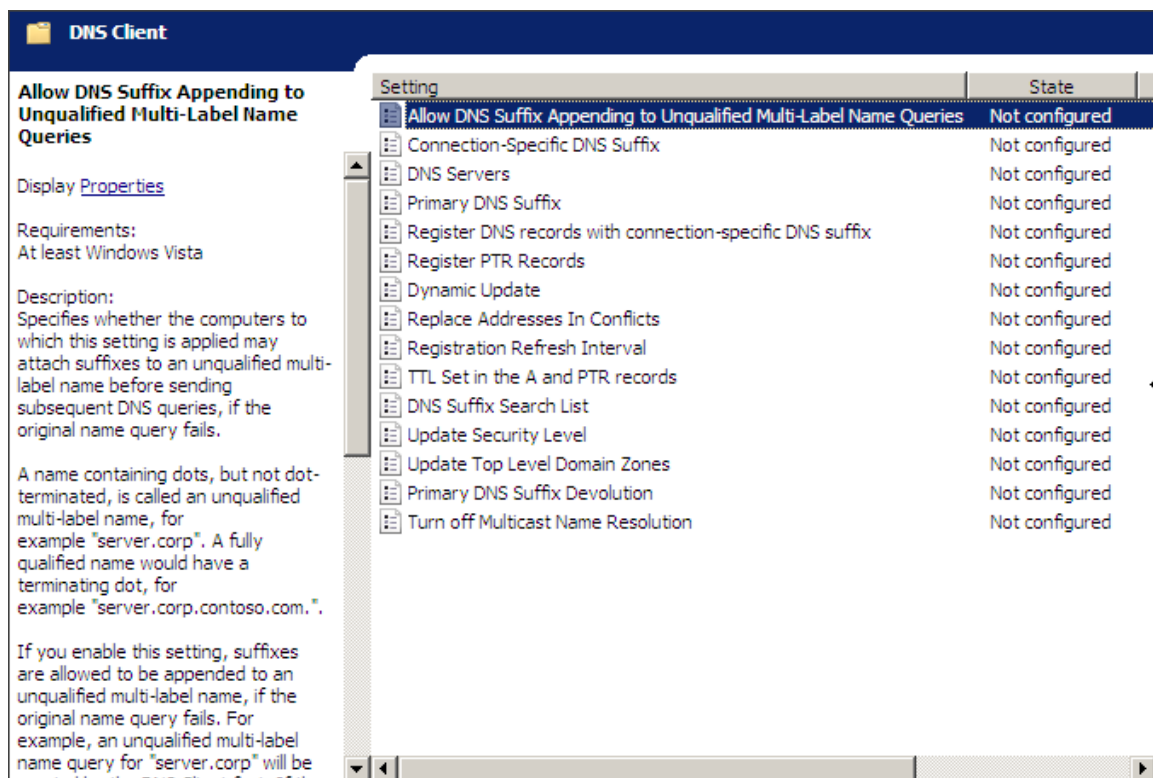


Figure 29: DNS Client Settings

Clearing the Resolver Cache

On the client side, the DNS resolver cache can become stale. This could cause the client to use outdated information for its name-resolution. To clear the resolve cache, do the following:

1. Open a **command prompt**.
2. Type **ipconfig /flushdns**.
3. Type **ipconfig /regdns**.
4. This will renew the DNS resolver cache.

DNS Server Lists

You can add additional DNS server lists to your client computer by doing the following:

1. Expand your **Network Card** in Network and Sharing Center (covered earlier).
2. Right-click IPv4 and choose **Properties**.
3. Click the **DNS** tab, which will bring up what you see below:

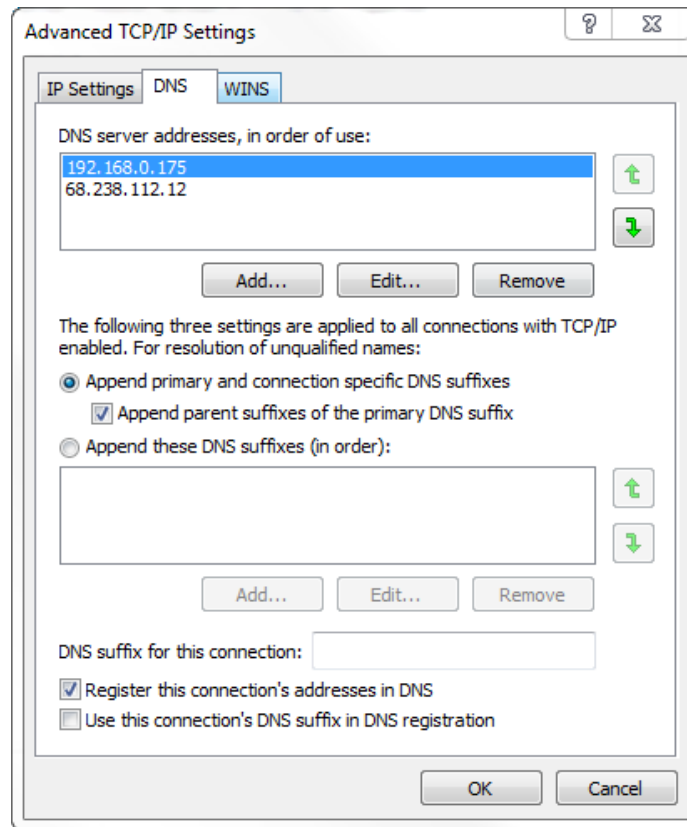


Figure 30: Client DNS Settings

4. Here, you can press the **Add** button and continue to enter the DNS address.

You can also change your DNS suffix in this location and change the order in which they appear:

1. Choose the **Append these DNS suffixes (in order):** radio button.
2. Click **Add**.
3. Here, you can choose the order of domain suffix. (For example, "preplonic.com.")

Link-Local Multicast Name Resolution

Traditionally, networks that don't have a DNS server have used NetBIOS to perform name resolution for other hosts on the network. However, NetBIOS isn't compatible with IPv6.

Link-local multicast name resolution, or LLMNR, is built into Windows Vista, Windows 7, and Windows Server 2008. It uses multi-casting to perform name-resolution on an IPv6 network in the absence of a DNS server. It's more secure than NetBIOS, but it can only be used on the local subnet. So if you need to enable single-label name resolution across a multi-subnet LAN, you would need to use DNS with GlobalNames zones instead.

To enable LLMNR, open the "Network and Sharing Center" and enable "Network Discovery."

Domain 3: Configuring Network Access

Configuring Network Access can involve one of the following four things:

1. Connecting two corporate networks together.
2. Allowing employees or customers to access the corporate network from an outside computer.
3. Allowing users to access the network via wireless technologies.
4. Allowing clients on the corporate network to access the Internet.

The Windows Server architecture has been designed to allow clients from all over the globe to easily connect to a Windows Server 2008 server. When Microsoft created these policies, the philosophy behind them was to create a centralized server that provided every feature that a client could possibly want in terms of remote access, all from a single platform. With Windows Server 2008, a single server is capable of supporting:

- Routing and Remote Access
- Network Access Protection
- Network Authentication
- Wireless Clients
- Application-Level Firewalls

In this section of the Mega Guide, we're going to explore all parts of network access. We'll begin by taking a look at a Network Address Translation.

Network Address Translation (NAT)

Since private IP addresses can't be routed across the Internet, we need a way for computers on a private LAN to communicate with computers on the Internet. We also need a way for a user at a remote location to connect a computer to the corporate LAN. The solution for both of these scenarios is to use Network Address Translation, or NAT. The NAT device would have a public IP address assigned to the Internet interface, and would direct Internet traffic to and from multiple LAN hosts via their private IP addresses.

NOTE: This only applies to IPv4. NAT isn't needed for IPv6.

Normally, you'd use a dedicated device, such as a router or a DSL gateway, to perform Network Address Translation. Since servers have to be booted more often and can be more prone to failure, dedicated NAT devices are usually more appropriate. However, for the exam, you still need to know how to set up a Windows Server 2008 machine as a NAT device. There are two ways to do that.

Internet Connection Sharing (ICS)

Internet Connection Sharing, or ICS, is extremely easy to set up. All it takes is a server with two network interfaces, and a few mouse clicks. It's more appropriate for small businesses, since its built-in DHCP server can only issue private IP addresses to clients on the local subnet. (And, you can't get around that problem by using a separate DHCP server. ICS isn't compatible with any DHCP server except for its own built-in implementation.) Here are the steps for setting it up:

1. Open **Control Panel**.
2. If **Network and Internet** is there, click it, otherwise click **Network and Sharing Center** and then click **Manage Network Connections**.
3. Right-click your connection and select **Properties**.
4. Click the **Sharing Tab**.

Routing and Remote Access Services (RRAS)

For a larger organization, it's more appropriate to set up NAT with Routing and Remote Access Services (RRAS). The advantage of using RRAS is that you can use it with any DHCP server. So you can use it on a LAN that's divided into multiple subnets.

The first step is to install the Network Policy and Access Services server role.

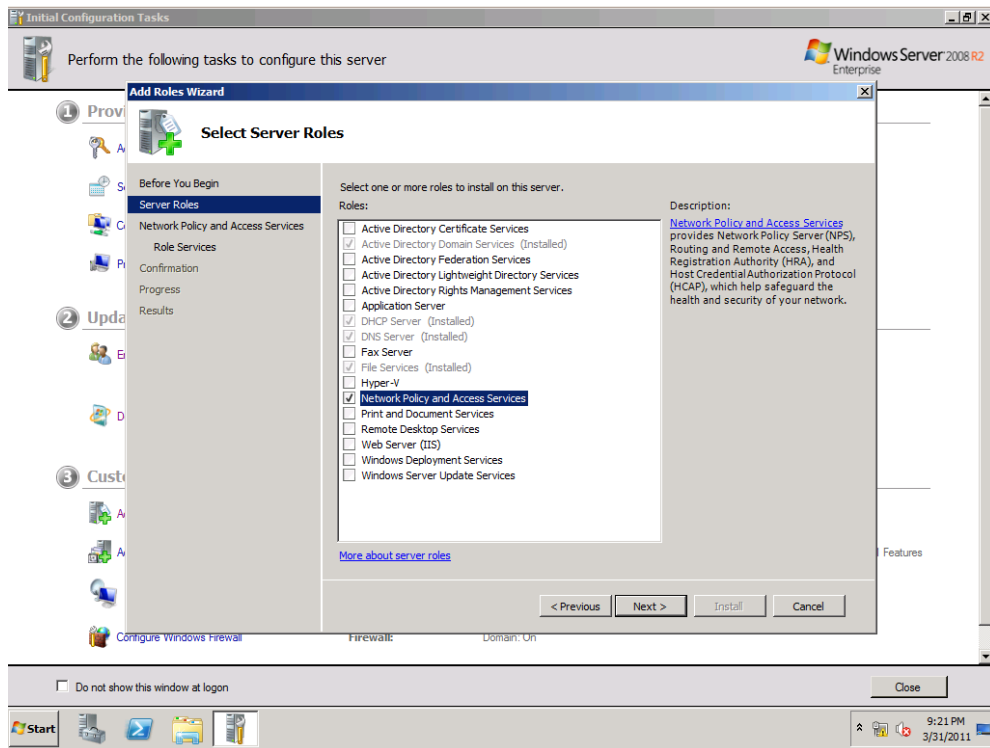


Figure 31: Selecting the Server Roles

On the Select Role Services page, select Routing and Remote Access Services.

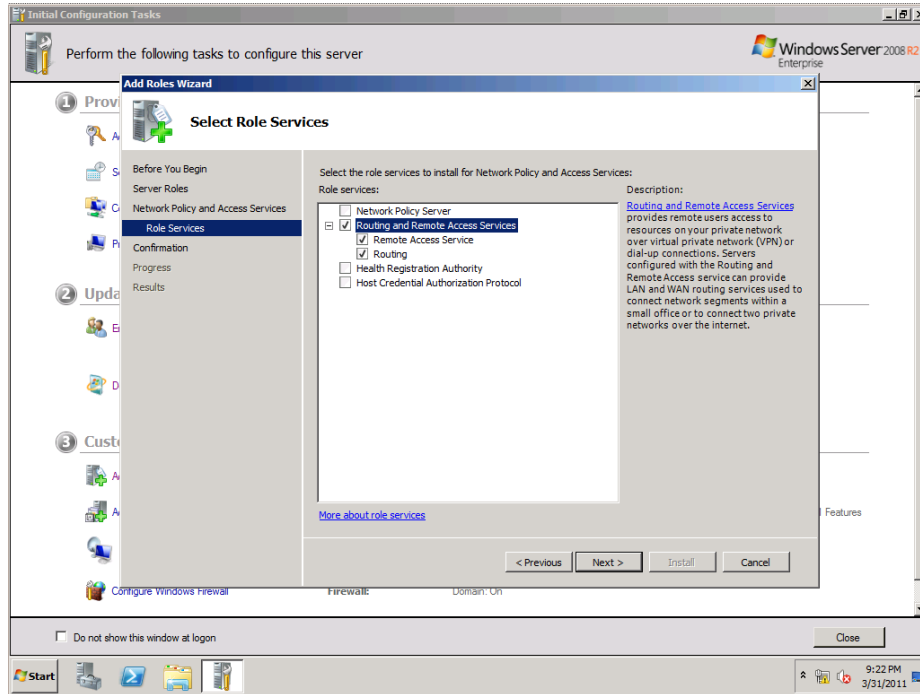


Figure 32: Installing RRAS

When installation is complete, open the RRAS snap-in to configure NAT.

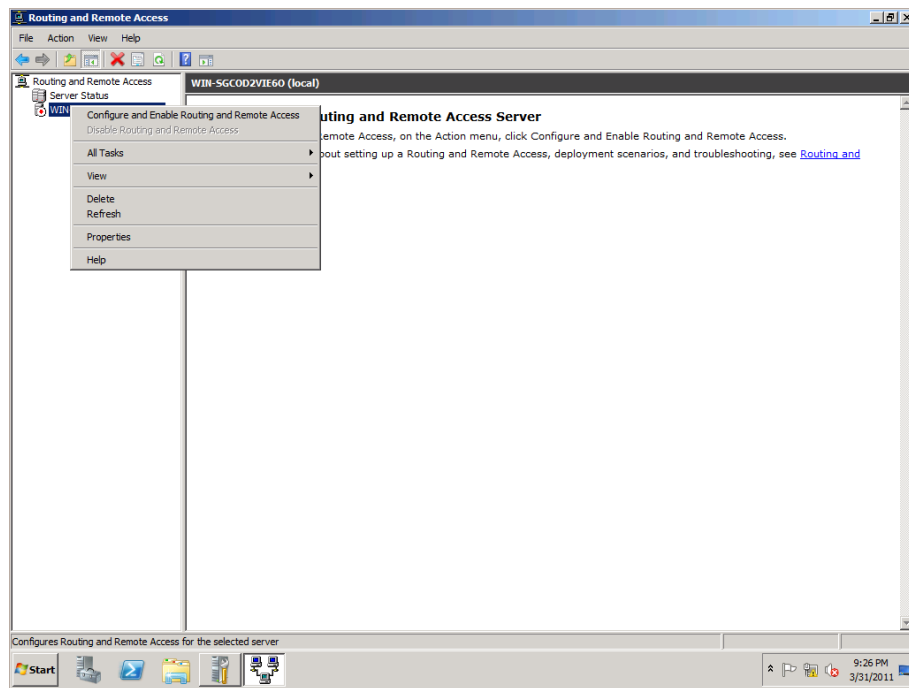


Figure 33: Configure RRAS

Now just follow through with the Wizard until setup is complete. Note that on the following screen, you'll only be able to use the demand-dial option unless your server has both a private network interface for the LAN, and a public network interface for the Internet.

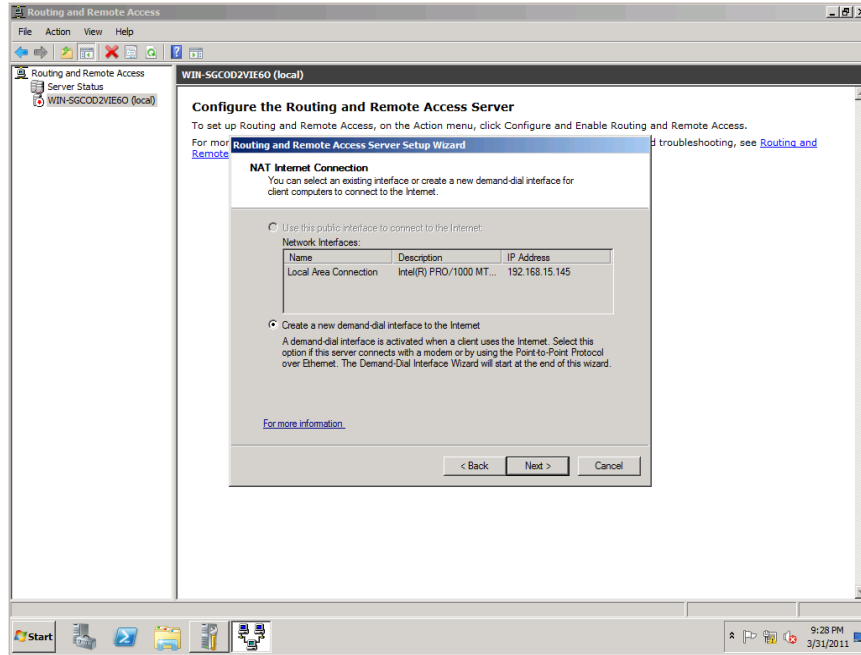


Figure 34: Configuring NAT

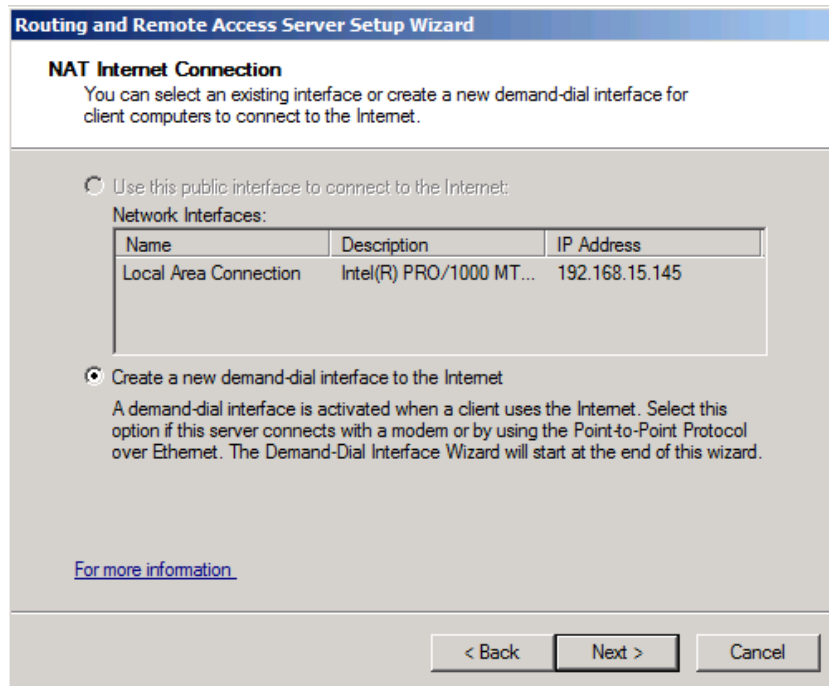


Figure 35: Configuring NAT

Inbound and Outbound Filters

You can set up inbound and outbound filters through the NAT portion of your Routing and Remote Access server by navigating to the Routing and Remote Access (R&RS) snap-in and expanding the IP routing section. Inbound and outbound filters monitor traffic that goes in and out of your network. For example, if someone has traffic going outbound, the outbound filters would apply. Inbound filters apply to inbound traffic.

Note: This does require you to have NAT set up already.

If you right-click on the NAT/Basic Firewall section, you'll be able to expand its properties. As you can see in the figure below, there will be an inbound and outbound filter setting. This will allow you to refine traffic based on your own criteria, including which addresses or protocols are allowed to access your network.

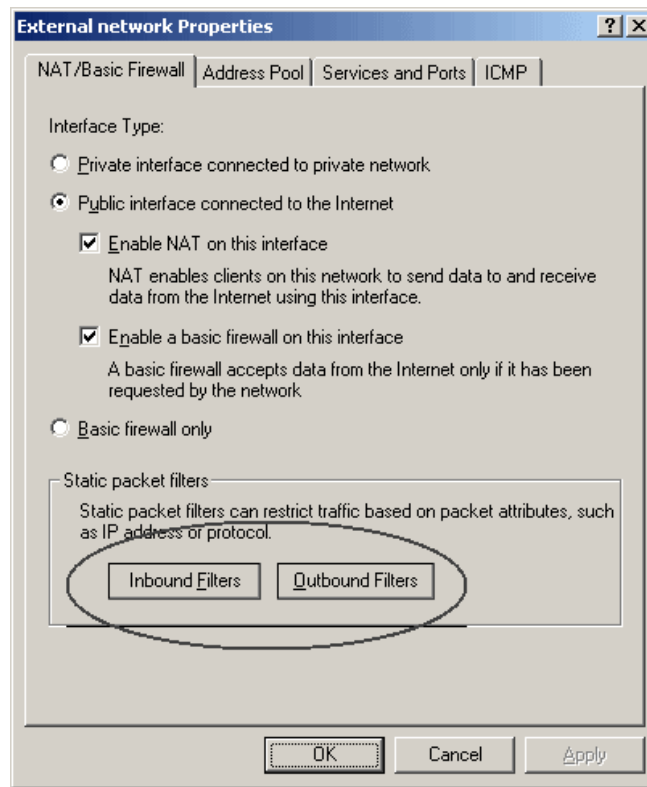


Figure 36: Creating Inbound and Outbound Filters

Configuring Remote Access

Once Network Address Translation has been set up, it's a simple matter to allow users on the LAN to access the Internet. However, your organization may need to allow some users to work from remote locations, such as from a customer's office or from home. Those remote workers may need to access resources on the corporate LAN. To allow this to happen, the organization will need to set up some method of remote access. With Windows Server 2008 R2, there are three ways to do this. You can use Demand Dial-Up via old-fashioned modem technology or use it with a Virtual Private Network. A third option is DirectAccess is a technology that's brand new for the R2 version of Windows Server 2008. Regardless of which access method is used, proper security measures must be implemented.

Remote Dial-Up Access

The first step is the same as for setting up Network Address Translation. That is, you'll need to install the Network Policy and Access Services server role, and from there install Routing and Remote Access Services.

To configure a dial-up connection, complete the following steps:

1. Navigate to **Start > Administrative Tools > Routing and Remote Access**.
2. Right-click on your server and select **Configure and Enable Remote Access**.
3. This will bring up the Routing and Remote Access Server Setup Wizard.
4. Select **Remote Access (Dial up or VPN)**.
5. Click **Next**.
6. Select the **Dial-Up** checkbox; click **Next**.
7. At the next screen, you can either choose a range of IPs specifically or assign them automatically.
8. Click **Next**.
9. You can choose to either have a RADIUS server perform authentication, or allow the destination server do it.
10. Click **Next**.
11. Click **Finish**.

Virtual Private Networks (VPN)

Although there may be times when users might still have to use a dial-up modem, it's becoming more common for remote users to connect to corporate resources with Virtual Private Networks, or VPNs. This allows users to access the corporate LAN in a secure manner across the Internet. Available bandwidth is much higher than it is with old-fashioned modems.

When setting up a VPN, you have the choice of either letting the connecting RRAS server perform its own user authentication, or of using a RADIUS server instead.

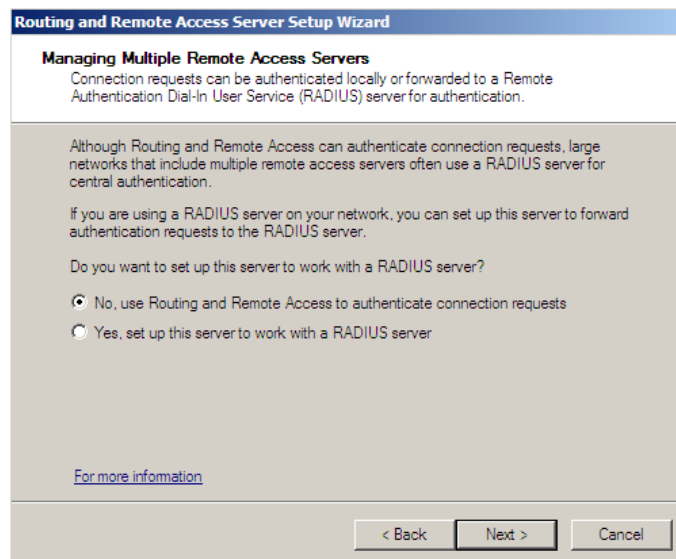


Figure 37: Configuring VPN Access for RADIUS

Windows Server 2008 offers three different protocols for use with a VPN.

PPTP

Point-to-Point Tunneling Protocol (PPTP) is the least complicated way to set up a VPN. It uses the Point-to-Point protocol (PPP) as its underlying transport mechanism. Generic Routing Encapsulation (GRE) is used to encapsulate PPP frames for transport. Because PPTP also uses normal PPP authentication, it is performed before the encrypted tunnel is set up. Encryption is performed with Microsoft Point-to-Point Encryption (MPPE) and the RC4 encryption algorithm.

Since PPTP doesn't use certificates, there's no need to set up a Certificate Authority. Also, PPTP can be used with Windows clients as far back as Windows 95.

PPTP provides data confidentiality, but it doesn't verify either data origin authenticity or data integrity.

L2TP/IPSec

This method is a combination of the Layer 2 Tunneling Protocol, or L2TP, and IPSec. L2TP is used to create the tunnel, and IPSec handles encryption and authentication. L2TP frames are encapsulated as User Datagram Protocol, or UDP, messages.

It's more complex to set up than PPTP, but it does offer a greater degree of security. Authentication is not performed until after the IPSec session is set up. It uses certificates for both the access servers and the clients, so a Certificate Authority will need to be set up. It can be used with clients as far back as Windows 2000.

L2TP/IPSec provides data confidentiality, verification of data origin authenticity, and verification of data integrity.

SSTP

Since both PPTP and L2TP/IPSec use their own ports, access could be blocked by a corporate firewall that may normally have those ports closed. The Secure Sockets Tunneling Protocol (SSTP) helps to get around that problem. It uses the same HTTP over SSL protocol that is used by secure web servers, so it will also use port 443, which is probably already open on the corporate firewall.

Although it uses certificates, it's still easier to administer than L2TP/IPSec. That's because SSTP only requires that certificates be issued to the access servers, and not to each individual client.

SSTP can be used with clients as far back as Windows XP SP3.

VPN Reconnect

Windows Server 2008 R2 can support VPN Reconnect when used with Windows 7 clients. It accomplishes this by using the IKEv2 Mobility and Multihoming protocol, or MOBIKE. It simply means that a user can move from one network connection to another, without losing the VPN. In this way, for example, a user could establish a VPN connection on a laptop that's hooked up to a wireless network in his office, unhook the network cable and carry the laptop to a conference room. There, he can connect to the wireless network, and the VPN connection will still be established. With other types of VPN access, the user would have to log back into the VPN every time the computer is moved to a different network.

DirectAccess

In reality, DirectAccess is just another type of VPN. However, Microsoft lists it as its own topic for the 70-642 exam. It's the most complex remote access method to set up, but it provides many benefits for remote users. Once it's set up, it can also make things easier for network administrators. It's brand new with Windows Server 2008 R2, and it can only be used with Windows 7 clients.

DirectAccess requires that IPv6 be operational on the LAN. It will then use a transitional technology, such as Teredo or 6to4, to enable access across the IPv4 Internet.

Although the remote access server must have Windows Server 2008 R2 installed, the other servers on the network only need to have the original Windows Server 2008. It also requires that a Certificate Authority and a high-availability website be set up on the LAN. Unlike traditional VPN servers, DirectAccess servers cannot be accessed behind a NAT device. A DirectAccess server must be set up with two network interfaces. One interface must be directly connected to the Internet and have two consecutive IPv4 addresses assigned to it. The other interface is connected to the LAN.

For the end user, the benefit is that operation is totally seamless. After logging on to the client computer, there's no need for the separate steps of creating a VPN connection and logging into the LAN.

When a user logs on to a computer, the computer first tries to contact the high-availability website on the company's Intranet. If the website can be contacted, then the computer knows that it's plugged directly into the LAN, and sets up normal LAN access. If the computer can't contact the website, then it knows that it's at a remote location, and automatically sets up the DirectAccess connection. Either way, after logon to the computer has been completed, the user has access to the corporate LAN.

For the system administrator, DirectAccess can greatly simplify user administration. Rather than having to create separate remote access policies for users, administrators can now just use the Group Policies that are already set up.

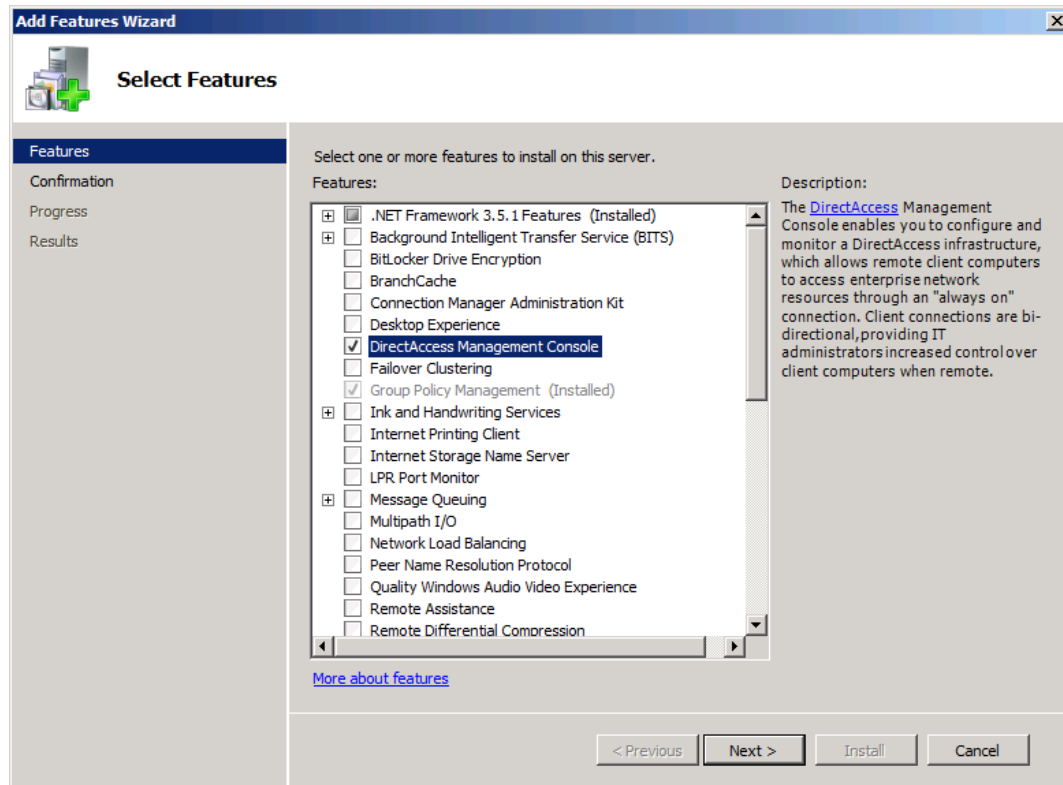


Figure 38: Use the Server Manager Wizard to Add the DirectAccess Console to Your RRAS Server.

Remote Authentication Protocols

A major part of remote access security involves ensuring that only authorized users can access the corporate LAN. It also involves ensuring that users who log in really are who they say they are, and ensuring that users' login credentials can't be captured by criminal hackers. Windows Server 2008 supports several authentication protocols.

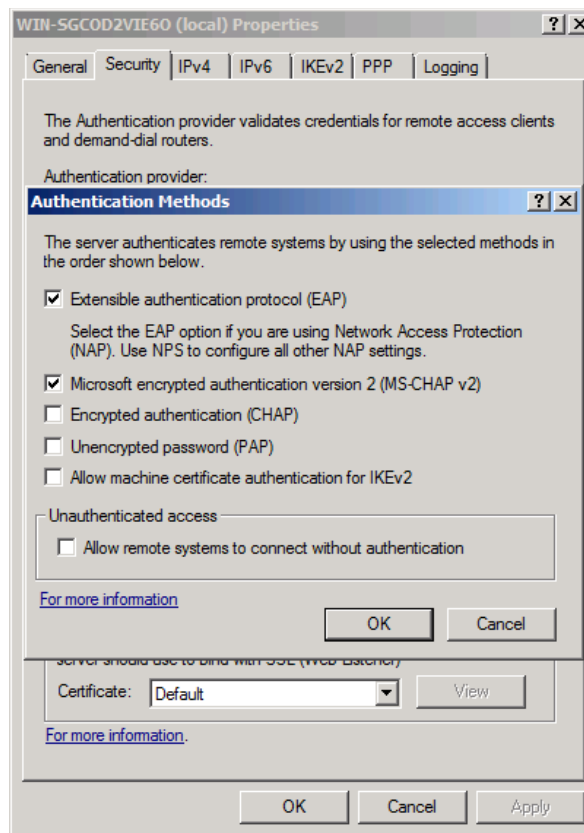


Figure 39: Authentication Methods Available on Windows Server 2008 R2

PAP

The Password Authentication Protocol sends users' login credentials to the server in plain-text. Although Windows Server 2008 R2 still supports it as an option, you really don't want to use it.

CHAP

The Challenge-Handshake Authentication Protocol, or CHAP, is an encrypted protocol which uses password authentication. However, the password itself is never sent across the network. Rather, when the client contacts the server, the server will send a challenge to the client. The client will then send a hashed value that it has calculated for both the challenge and the user's password. The server will calculate its own hashed values, and compare them with what it receives from the client. At random intervals, the server will send challenges to the client, repeating the authentication process.

MS-CHAPv2

MS-CHAPv2 is the Microsoft implementation of CHAP. Its biggest enhancement over regular CHAP is its ability to perform mutual authentication. In other words, instead of just having the client authenticate to the server, the server will also authenticate to the client. This way, users can be sure that they're not logging into a counterfeit server.

EAP and PEAP

The Extensible Authentication Protocol (EAP) and the Protected Extensible Authentication Protocol (PEAP) both use either certificates or smart cards. Either way, a certificate infrastructure must be set up, which makes EAP and PEAP more complex to implement. However, they make for better security than either CHAP or MS-CHAP.

The difference between EAP and PEAP is that PEAP creates an encrypted channel for EAP communications. It does this by setting up a Transport Layer Security, or TLS, tunnel.

EAP and PEAP are widely used for wireless access. PEAP can be used along with MS-CHAPv2 to enable two-factor authentication. (In other words, both a password and a certificate or smart card are used, instead of just one or the other.)

DHCP Considerations for Mobile Clients

One thing you'll want to watch out for when configuring mobile clients is the lease duration for their IP addresses. If you allow long lease durations for mobile clients, then their IP addresses will remain in effect even when the clients have been removed from the network. This could result in a shortage of IP addresses that may be issued by your DHCP server. Therefore, your best bet is to configure the DHCP server to issue short-duration leases on IP addresses for mobile clients.

RADIUS and Network Policy Server

User authentication can be performed by either the Remote Access Server, or by a RADIUS server. RADIUS is an "authentication, authorization, and accounting" protocol.

When a Remote Access Server receives a connection request from a client, it will pass the request along to the RADIUS server. The RADIUS server will then authenticate the user and pass that information to the Remote Access Server. It will also inform the Remote Access Server about what resources the user is allowed to access, and will keep a log of network connections, along with data about the users who make the connections.

Older implementations of RADIUS maintained user credential data in a flat-file database. Newer RADIUS implementations can use either the flat-file database, an SQL database, LDAP, or Active Directory.

RADIUS can be used to authenticate users who connect via a VPN, wireless, or even old-fashioned dial-up modem.

Note that Microsoft has built its Windows Server 2008 implementation of RADIUS into the Network Policy Server. (Network Policy Server has replaced the Internet Authentication Service that came with Windows Server 2003.)

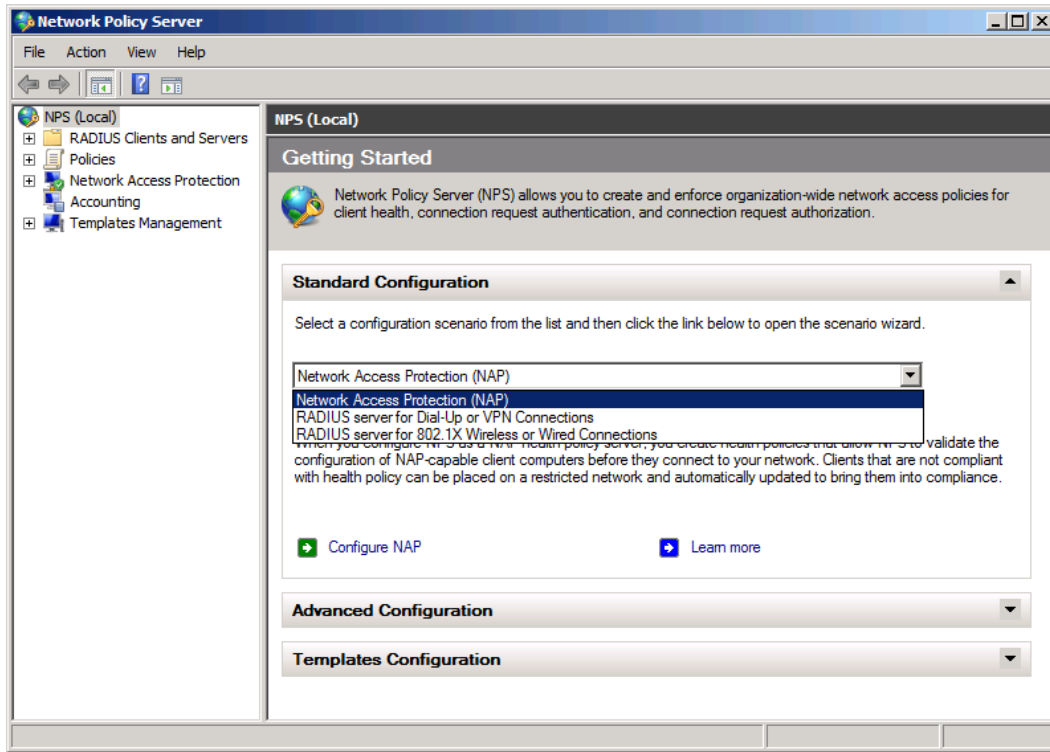


Figure 40: Setting Up a RADIUS Server

RADIUS Clients, Server, and Proxies

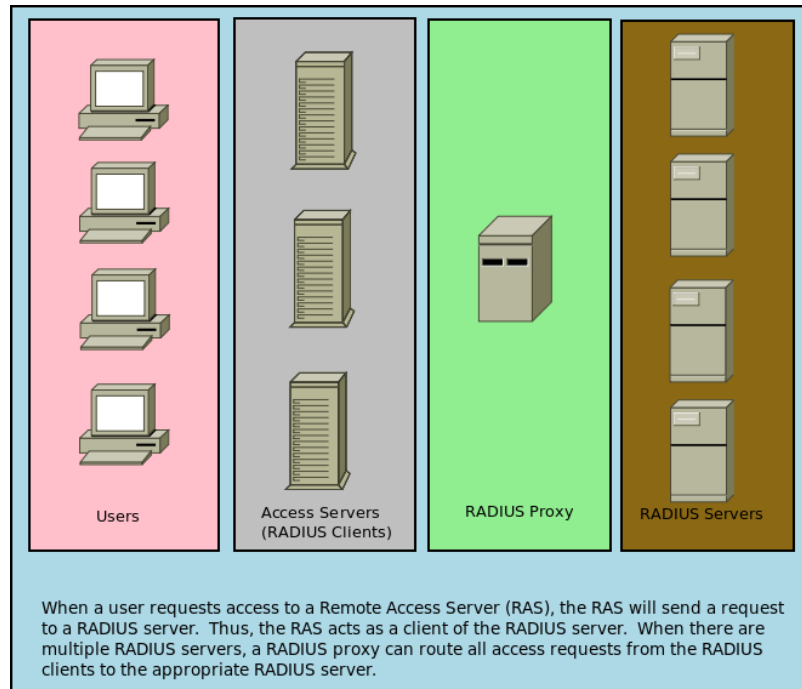


Figure 41: RADIUS Client/Proxy/Server Relationship

Configure Network Access Protection (NAP)

Network Access Protection is designed to check the health of any client computers that try to join the network. It checks the clients for the presence of up-to-date anti-virus programs, anti-spyware programs, and a host-based firewall. NAP also verifies that the most recent Windows updates have been installed.

When a host computer tries to join the network, it will submit its health status to an enforcement point. (An enforcement point can be a DHCP server, an IEEE 802.1x switch, an IEEE 802.11 wireless access point, a VPN server, or a Health Registration Authority.) The enforcement point will use the RADIUS protocol to send the host's health status to an NAP health policy server for evaluation. (Note that this same NAP server can also provide RADIUS authentication services.)

If the NAP server determines that the host meets the proper health criteria, then the host will be allowed to join the network. Hosts that do not meet the health criteria will normally either be denied access, or be directed to a quarantine network for remediation. You also have the choice of allowing the unhealthy host on the network regardless, only using NAP for accounting purposes.

Although NAP can be used with the corporate desktop computers, it's especially useful for mobile users, visitors who bring in their own laptops, and employees who work from home with their own computers. In all of these situations, NAP can prevent an unhealthy computer from joining the corporate network.

DHCP Enforcement

By using the NAP configuration Wizard, you can configure DHCP to enforce NAP policies. Before issuing or renewing an IP address lease, the DHCP server will submit the client's health report to the health policy server. If the health policy server tells the DHCP server that the client does not meet health criteria, the DHCP server can prevent the client from joining the network. The DHCP server can either deny access altogether, or issue the IP address of the restricted, remediation network.

DHCP enforcement is the least secure of all NAP enforcement methods. It won't work with clients that are configured with static IP addresses. Also, any user who has administrative privileges on the client computer can override NAP simply by changing the client to a static IP address.

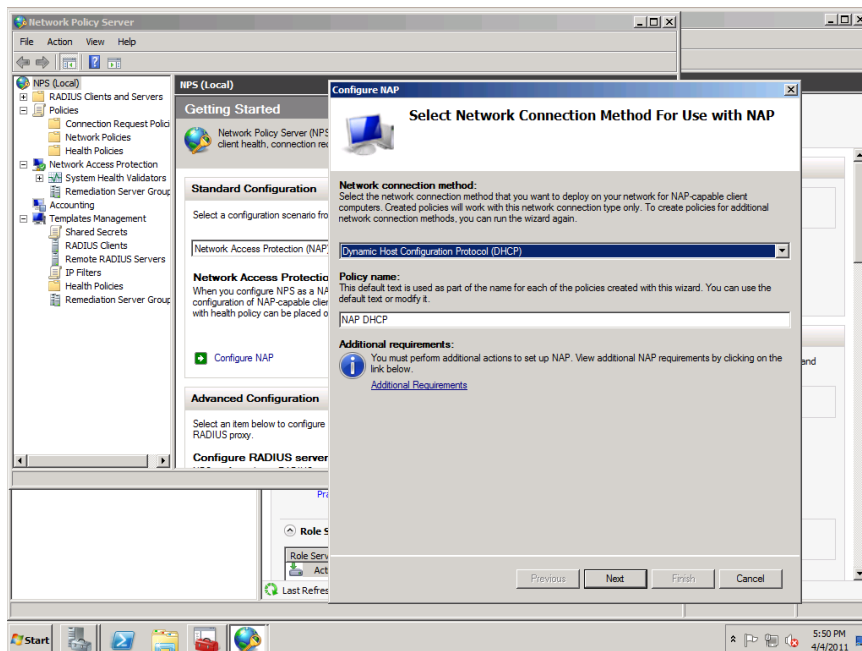


Figure 42: Configuring NAP

VPN enforcement

VPN enforcement simply means that only healthy clients are allowed access to a Virtual Private Network. It only works with clients that run Windows XP SP3 or higher.

With VPN enforcement, the client's health is constantly monitored. That way, if a client that was considered healthy when it first logged into the VPN suddenly becomes unhealthy, it can be taken out of the main network and placed into the restricted, remedial network.

VPN enforcement is a very strong type of enforcement. It requires the use of certificates, which means that a Certificate Authority will have to be set up on the network.

IPSec with Health Registration Authority Enforcement

With this type of enforcement, computers must meet health criteria before they can begin IPSec communications with other computers. By leveraging the power of IPSec rules, health requirements can be customized on an IP address or communications port basis.

This is the strongest type of NAP enforcement. It works with client operating systems as far back as Windows XP SP3.

802.1x Enforcement

802.1X is an implementation method that takes advantage of 802.1X compatible hardware, such as switches or wireless access points. Non-healthy computers that try to join the network are directed to the restricted, remedial network.

The health status of clients is constantly monitored, and clients that become unhealthy are shifted over to the restricted, remedial network. This type of enforcement works with clients that are running Windows XP SP3 and newer operating systems.

This is considered a strong type of NAP enforcement. A Certificate Authority infrastructure is required.

System Health Agents and System Health Validators

A System Health Agent, or SHA, is installed on each client that is to be monitored. Windows XP SP3, Windows Vista, and Windows 7 all come with an SHA that will monitor the Windows Security Center. The clients' SHAs can be configured via Group Policy.

The System Health Validator (SHV) is installed on a Windows Server 2008 machine. This is what processes the information that comes from each client's SHA. A new feature of Windows Server 2008 R2 is the ability to set up multiple SHVs on the same set of NAP health policy servers. This allows administrators to have multiple sets of health policies for different purposes.

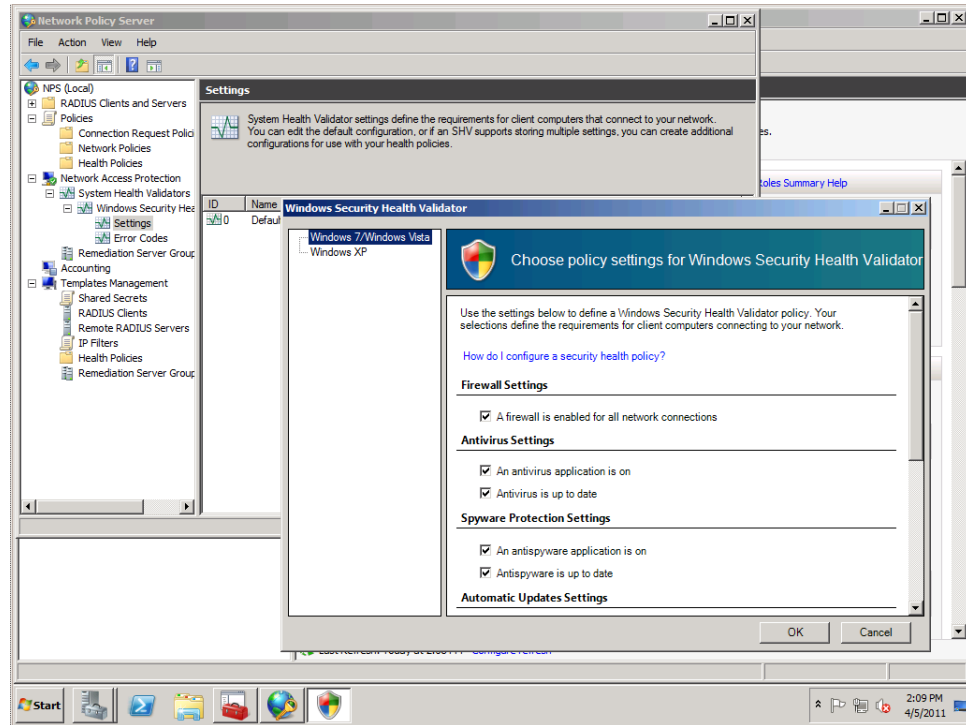


Figure 43: A Health Validator

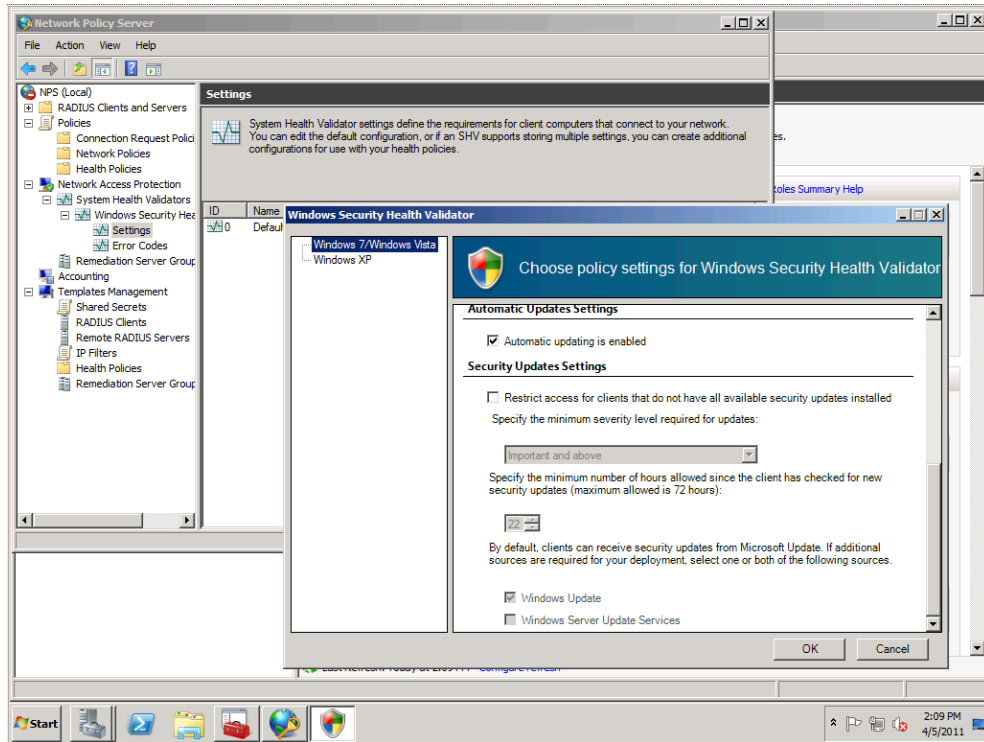


Figure 44: Configuring Health Validator Compliance

Note that NAP can be configured in Group Policy, using the Group Policy Management snap-in.

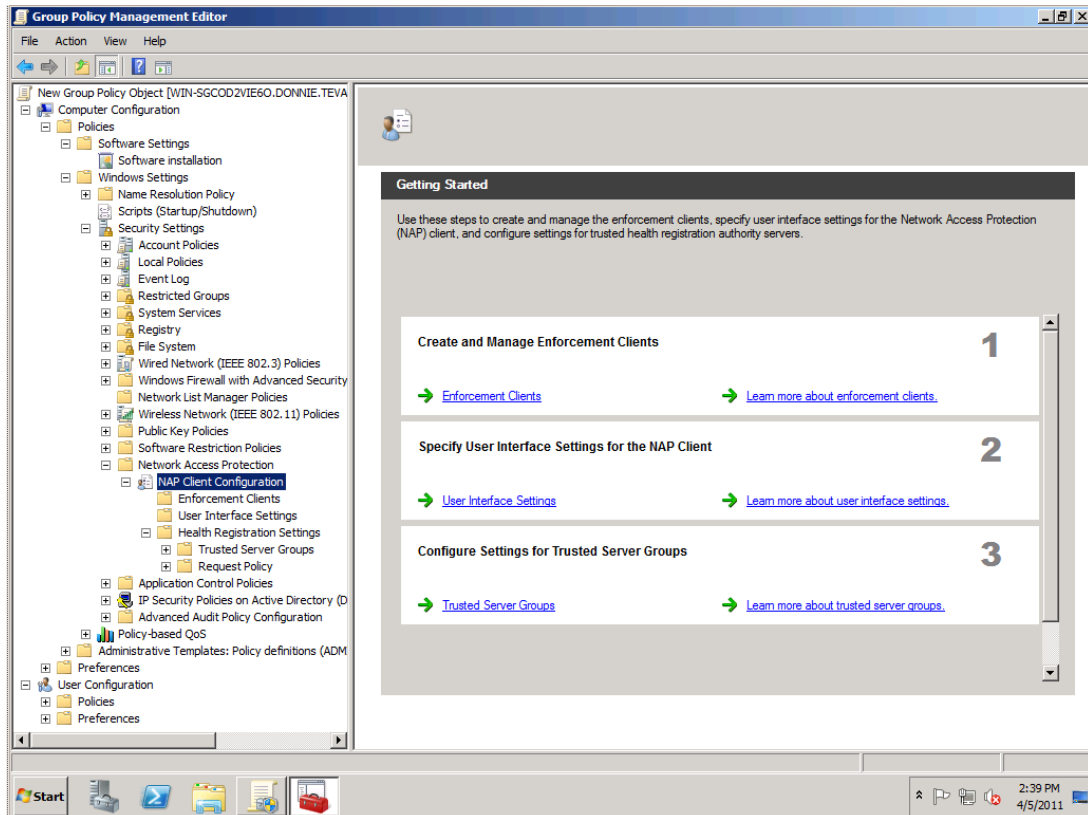


Figure 45: Configuring Group Policy for NAP

Configuring Wireless Access

Wireless access used to be found only in places like coffee shops and hotel lobbies, so that customers could go in and have easy Internet access. Now, though, wireless is becoming more prevalent in the corporate workplace. Among its key advantages are its portable nature, and the fact that its physical infrastructure is easier to install than that of a wired network.

However, along with the advantages is one disadvantage. That is, it's harder to make a wireless network secure than it is for a wired network.

Encryption Technologies for Wireless

No Encryption

This is what you encounter when you take your laptop to a public place that offers wi-fi access. Everything goes across in the open, with nothing to stop snoopers from intercepting your communications. Of course, this isn't what you want when exchanging sensitive data.

If you absolutely have to use a public, non-encrypted wireless access point, be aware of the dangers. Also, using SSL encryption, such as when accessing a secure website, can help somewhat.

Wired Equivalent Privacy (WEP)

Other than having no encryption at all, WEP is the most insecure of all the wireless encryption technologies. It uses a hexadecimal pre-shared key to encrypt the signal on both the client and the server ends. WEP can use 40-, 64-, and 128-bit encryption keys. However, no matter the key size, it is considered highly insecure because it is easily crackable. It should never be used when the exchange of sensitive data is involved. You may find some devices, like wireless printers, that support nothing but WEP.

Wi-Fi Protected Access with Pre-Shared Key (WPA-PSK)

In general, WPA is a much better encryption technology than WEP. However, even WPA's effectiveness can vary, depending upon how you deploy it.

WPA-PSK uses a static, pre-shared key, the same as WEP does. This means that a criminal hacker could conceivably use brute force methods to crack the key and enter the network. Also, it's very difficult to manage a set of static keys in a large corporate environment. If even just one computer with a wireless key is compromised, then the key would have to be changed on every wireless access point.

WPA-PSK is also known as WPA-Personal.

Wi-Fi Protected Access with Extensible Authentication Protocol (WPA-EAP)

WPA-EAP uses certificates and RADIUS for authentication. Optionally, it can also be used with smart cards. This is much more secure than WPA-PSK. It's also known as WPA-Enterprise.

WPA2-PSK and WPA2-EAP

WPA2 isn't an updated version of WPA. Rather, it's a whole new technology that's even more secure than WPA. It also comes in PSK and EAP varieties.

Of all the encryption technologies we've covered, WPA2-EAP is the strongest. It's recommended for use whenever possible.

Wireless Protocols

Wireless networking protocols are defined by the IEEE 802.11x family of specifications. You need to be familiar with them for the exam.

802.11a is an obsolete standard that's no longer used.

802.11b is the most commonly used wireless protocol. It has an advertised speed of 11 Mbps.

802.11g is an update to 802.11b. Its advertised speed is 54 Mbps. You can use this protocol in mixed mode, so that an access point will also support clients that still run 802.11b. However, this will cause everyone's throughput to be lowered to the 802.11b speed. You can also run this in 802.11g-only mode, which means that access points will only support clients that are running 802.11g. This also means that everyone can take advantage of 802.11g's higher throughput speeds.

802.11n is an update to 802.11g. It provides an improved range, and an advertised throughput speed of 250 Mbps.

NOTE: You do need to be familiar with the advertised throughput speeds of the above wireless protocols. However, realize that in real life, the actual throughput speeds may be much lower.

Ad Hoc vs. Infrastructure Mode

An Ad Hoc wireless network is just when two or more computers communicate directly with each other via wireless technology. No wireless access point is involved. Ad hoc networks are useful for just transferring data between two computers “over the air.” You might use this if you had two wireless laptops and wanted to get data from one place to another, without using any additional hardware. However, there is a danger involved. Ad hoc networks make it very easy for criminal hackers to gain access to your machine.

In infrastructure mode, computers can only access the network via a wireless access point. You can think of the wireless access point as a central hub, much like an Ethernet switch on a wired network. Most wireless networks in enterprise environments are set up in infrastructure mode.

Group Policy for Wireless

For the exam, you should know that Group Policy does support wireless networks and that you should place your group policy clients into a dedicated Organizational Unit. This way, you can apply GPOs to them in an isolated manner.

Note that there are two general types of Group Policies that you can create for wireless clients. You can create policies for either Windows XP clients or for Windows Vista/7 clients. The following screenshot shows the policies that can be set for Windows Vista and Windows 7 clients. Note the bottom of the screen, where there are three settings that are specific to Windows 7.

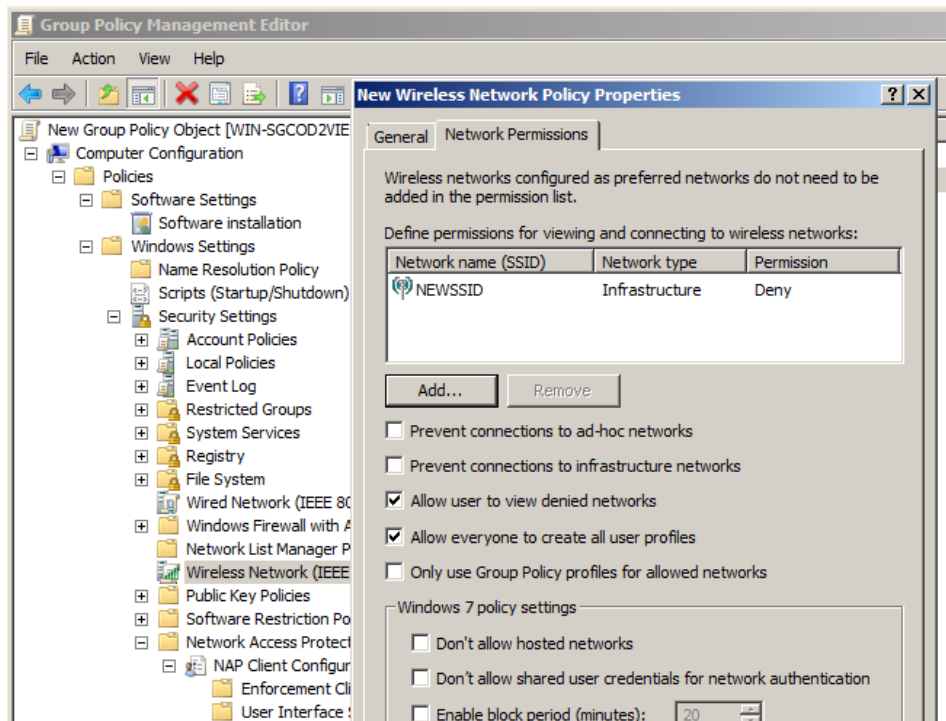


Figure 46: Configuring Group Policy for Windows Vista and Windows 7 Wireless Clients

Realize that when you configure Group Policies for wireless laptops, you’ll need to plug the laptops into the wired network, log on to the network, and log off. Otherwise, the users will never be able to log on to the wireless network, since the wireless Group Policies were never transferred to the laptop.

Configuring Firewall Settings

Windows Server 2008 contains a very modular firewall that can filter both incoming and outgoing traffic based on particular rules. With Windows Server 2008, you can use either the standard Windows firewall, or the new Windows Firewall with Advanced Security. You can either configure firewall settings for the local computer, or configure a Group Policy that will apply the same firewall settings to a group of computers.

Creating incoming or outgoing rules in Windows Server 2008 is a very simple process. First, select **Inbound or Outbound** in the Windows Server firewall properties. Then select **New Rule**. In the Wizard, you can refine your policy based on several fields, including port, protocol, and application. Make sure you practice setting up filters in your home lab.

Firewall Profiles

The Windows Firewall supports the ability to apply Firewall rules based on where a computer is located. The rules for the different locations are grouped into three different firewall profiles. They are:

- **Domain** – A domain profile is used when computers are joined and authenticated to an Active Directory domain. This profile will be applied any time a domain member's domain controller can be accessed.
- **Private** – By default, no networks are designated as private. To use this profile, a user will have to manually designate a network, such as a home network or small office network, as private.
- **Public** – By default, this profile will apply any time that a domain controller is not accessible. The default settings for this profile allow all outgoing traffic, but block all incoming traffic that isn't in reply to an outgoing connection request.

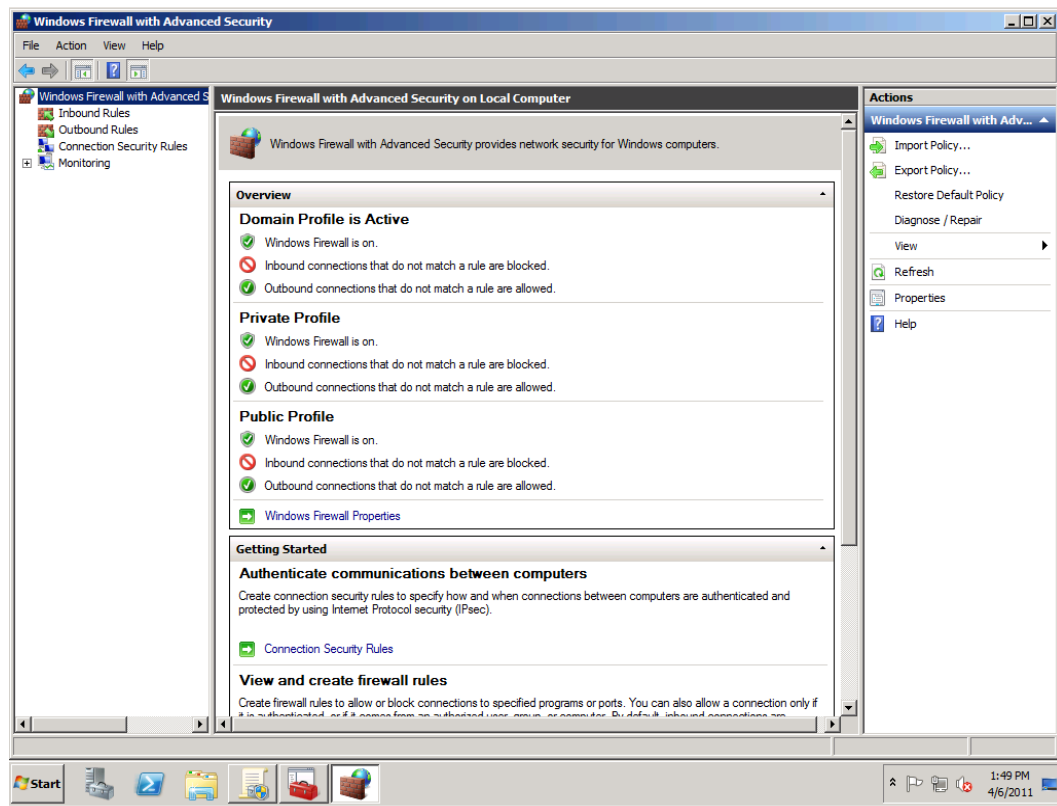


Figure 47: Windows Firewall Profiles

By default, all of these profiles block incoming traffic that isn't specifically allowed. On the other hand, all of these profiles allow all outgoing traffic. If what's in these profiles isn't satisfactory for your organization, you can customize rules for either incoming or outgoing traffic.

There are several reasons you may want to create rules for outgoing traffic. If one computer on your network gets infected with a worm, blocking certain outgoing firewall ports can help prevent the worm from spreading to other computers. Also, you may want to block certain outgoing ports to prevent employees from using unapproved applications that may send sensitive data out onto the Internet.

Note that when you install Microsoft applications on a computer, the firewall will normally be automatically configured to allow those applications to function. However, the installation programs for third-party applications may not configure the firewall for you. For this reason, when installing third-party applications, you may have to manually tweak the firewall rules so that the application can function.




Figure 48: Setting Rules for Windows Firewall with Advanced Security

Identify Ports and Protocols

You should be familiar with the following ports and protocols:

Note: Protocols are TCP unless otherwise noted.

Protocol Name	Port Number	Description
File Transfer Protocol	20	UDP protocol used to transmit Files
File Transfer Protocol	21	TCP version of FTP used to transmit files more quickly and accurately
SSH (Secure Shell)	22	Protocol used to establish command shells over distances securely
Telnet	23	Used to establish insecure shells over distances
SMTP (Send Mail Transfer Protocol)	25	Used to transport mail
whois	43	Used to lookup IP/name information
DNS (Domain Name Service)	53	Used to support transmission of IP/Name conversion information
DHCP (Dynamic Host Control Protocol)	68	Used to automatically obtain IP addresses
HTTP (HyperText Transfer Protocol)	80	Used to transmit data over the world wide web / Internet
POP3 (Post Office Protocol, version 3)	110	Used to receive email from an email server
SFTP (Secure File Transfer Protocol)	115	Security version of FTP
NTP (Network Time Protocol)	123	UDP protocol used to transmit network time
IMAP (Internet Message Access Protocol)	143	Protocol used to access email from a server, similar to POP3
SNMP (Simple Network Management Protocol)	161	UDP protocol used to monitor network attached devices
LDAP (Lightweight Directory Access Protocol)	389	Used to query active directory information via a network, even with Non-Windows based platforms
L2TP/IPSec	500 and 1701	Port 500 UDP used for IPSec negotiation Port 1701 used for the L2TP tunnel
SSL (Secure Socket Layer) This is also used for SSTP-type VPNs	443	Security protocol used to exchange SSL certificates
PPTP (Point-to-Point Tunneling Protocol)	1723	TCP and UDP. For VPN traffic
IKEv2 (MOBIKE)	4500	For VPN Reconnect
Remote Desktop Protocol	3389	For Remote Desktop Protocol traffic

Figure 49: TCP and UDP Ports

Firewall Scope

You can set up a firewall scope that will allow connections from the internal LAN while blocking connections from outside of the LAN. There are four different ways that you can set up a firewall scope.

- If you have a server that's connected to the Internet, such as a web server, the firewall scope can be configured to allow public access. At the same time, the scope can also prevent users of the internal LAN from accessing anything except for internal servers.
- Scopes for internal servers can be configured to allow access from only specified subnets. However, when planning for this, remember to include any remote access subnets that may be required by mobile users.
- For outgoing connections from internal servers, you can configure a scope that will allow specific applications to only connect to other specific internal servers.
- Scopes for mobile users can be set to allow traffic from specified services, such as Remote Desktop, to only communicate on certain subnets.

Firewall scopes can be configured in the Windows Firewall with Advanced Security snap-in.

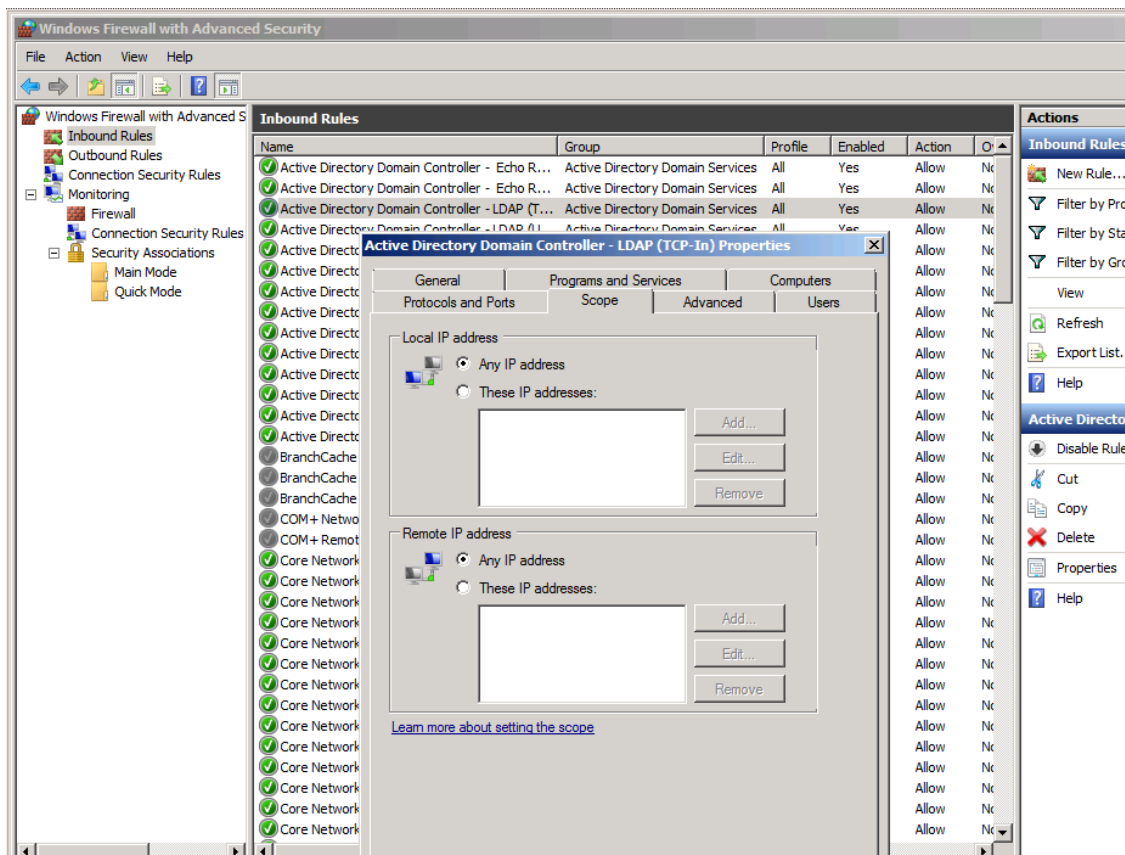


Figure 50: Setting a Firewall Scope

Configuring the Firewall through Group Policy

You can configure the Windows Firewall through group policy by opening the Group Policy Management Editor and navigating to **Computer Configuration > Policies > Windows Settings > Windows Firewall with Advanced Security**. There you can modify the policy just as you would with the normal Windows Firewall. This policy can then be applied to groups of computers in the domain.

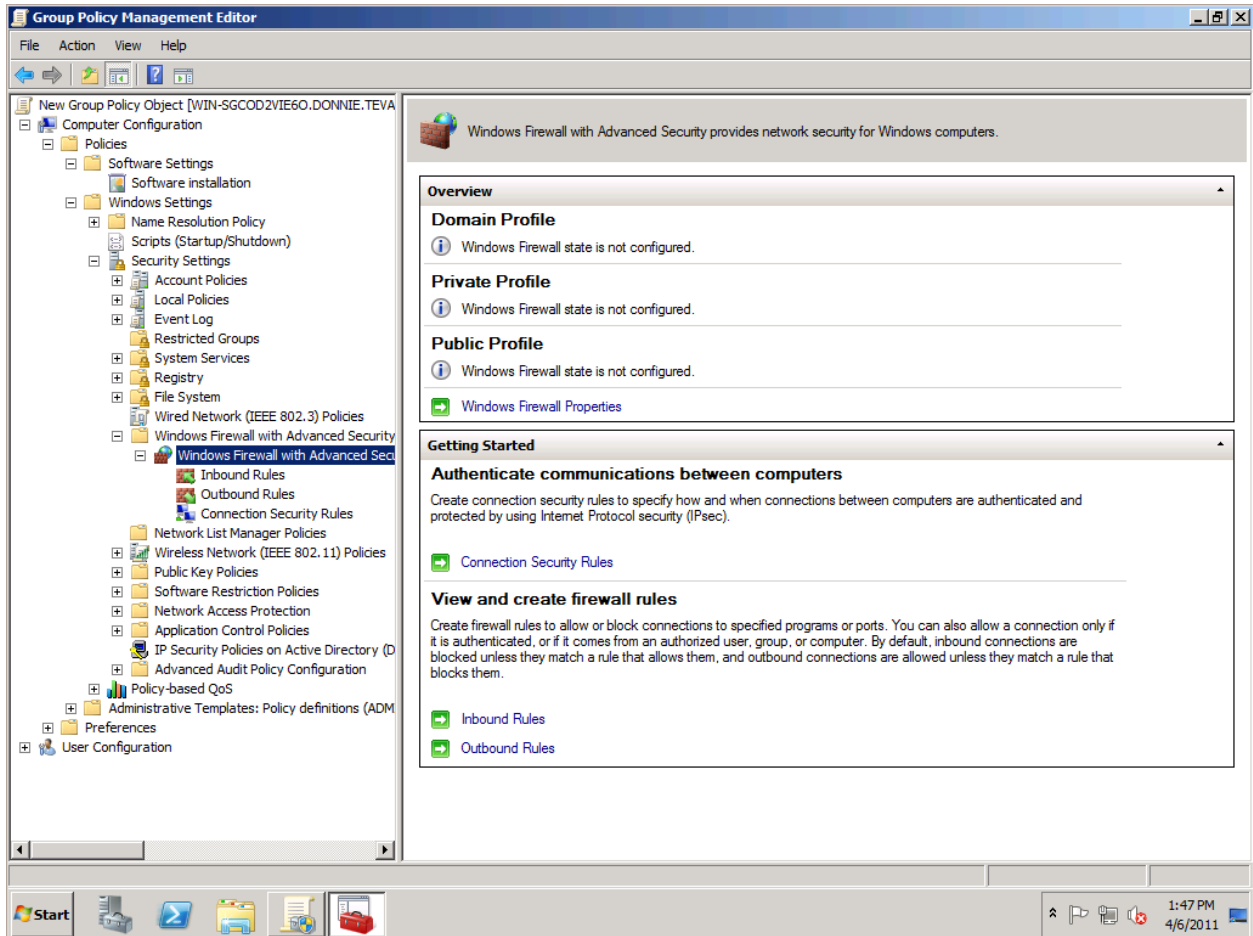


Figure 51: Configuring Firewall Rules in Group Policy

Authorizing Connections to Applications without Built-In Access Control

Many applications, such as the IIS web server that comes with Windows Server 2008, have their own built-in access control mechanisms. For example, you can set up a web server so that only certain authorized users can log in to a specified web page.

Other applications, however, may not have their own access control mechanisms. For them, you can use the advanced Windows Firewall along with IPSec to easily create your own access control.

For example, let's say you have your own customized finance application that listens on server port 1172. You can use IPSec to allow connections from only authorized computers or authorized users to connect to that port.

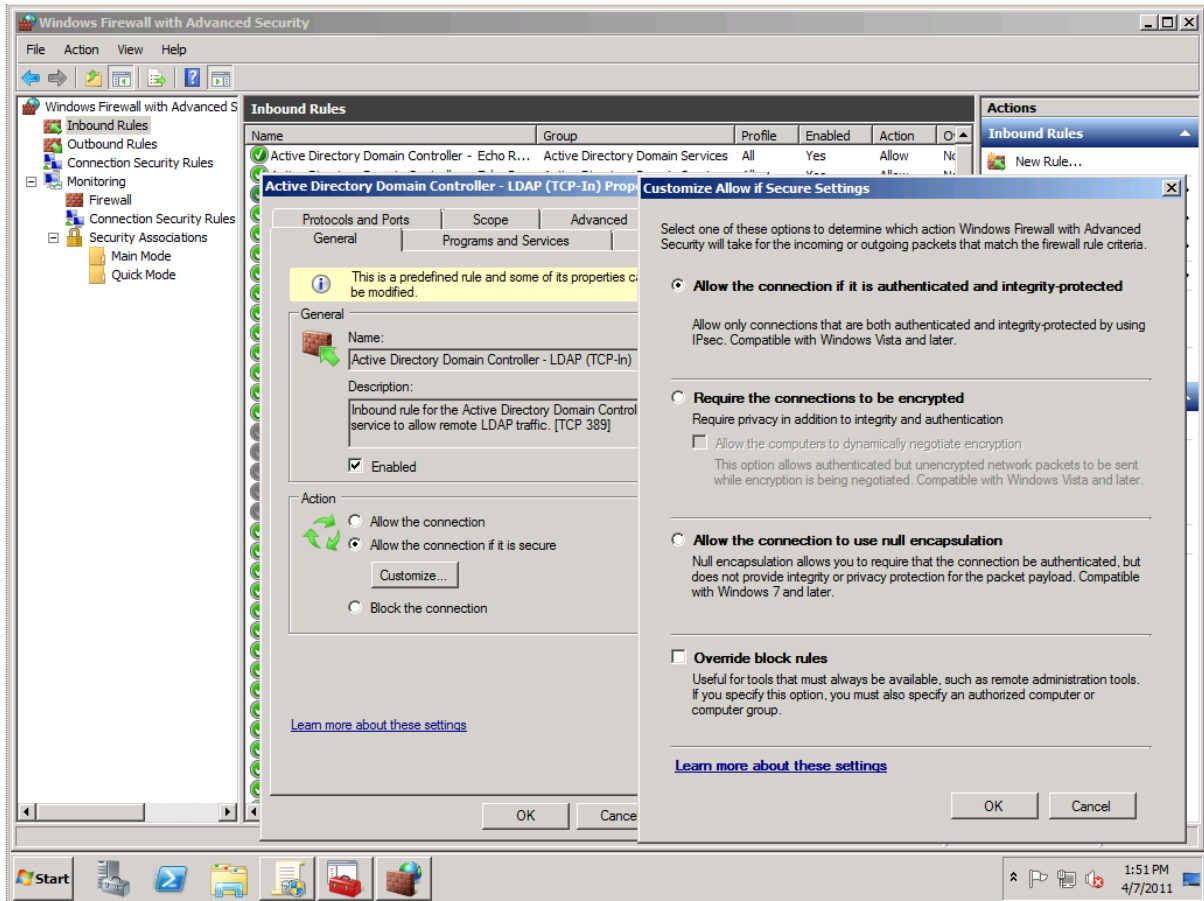


Figure 52: Configuring Access Control with Windows Firewall

Domain 4: Configuring File and Print Services

Windows Server 2008 supports the ability to share files, folders, and printers throughout your network using Active Directory Domain Services and simple network protocols.

Configuring a File Server

It's not necessary to add any new server roles in order to share files and folders; you can simply use the Windows Explorer. However, by adding the File Services role, you'll be able to make use of the Share and Storage Management snap-in.

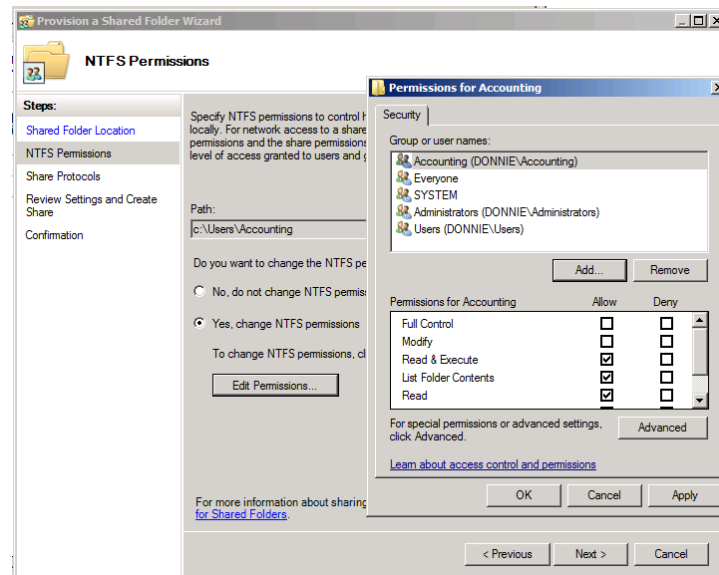


Figure 53: Share and Storage Management Snap-In

File Server Resource Manager

Another benefit of adding the File Services role is that you'll also be able to install the File Server Resource Manager. This is a set of disk and file management tools that was introduced in Windows Server 2003 R2. To install it, first install the File Services role. Then use the Server Manager to install the File Server Resource Manager.

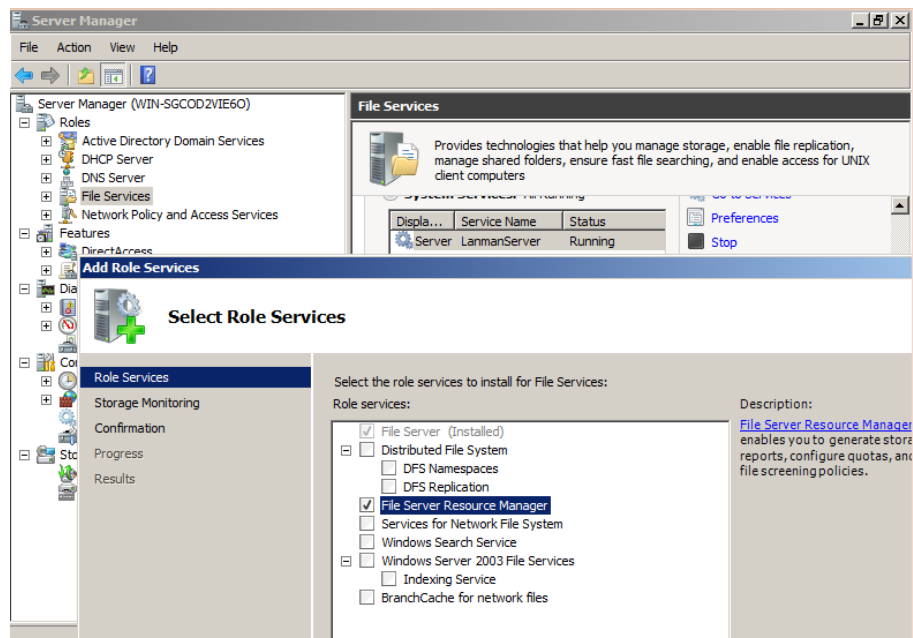


Figure 54: Installing the File Server Resource Manager

There are three main functions that you can perform from within the File Server Resource Manager.

Quota Management

File Server Resource Manager gives you an easy way to manage disk usage quotas for users. You can set usage quotas for either an entire volume, or just on specific folders. There's even a set of pre-defined templates to help you get started. If a user comes close to exceeding his or her quota, File Server Resource Manager can send an alert to the appropriate administrators.

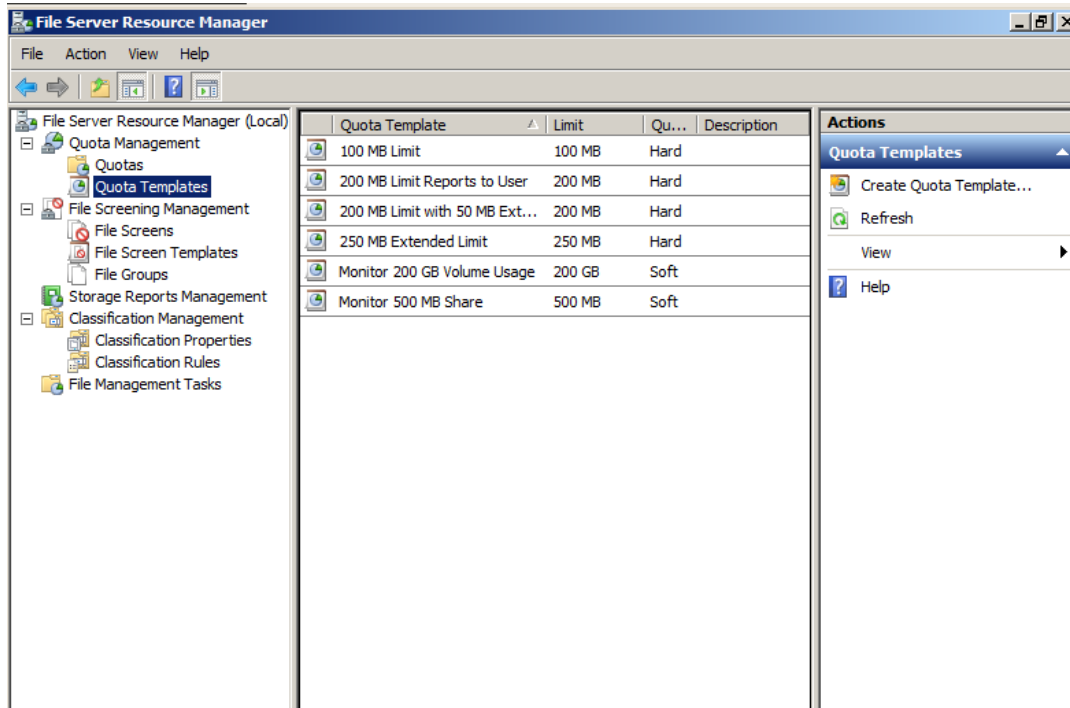


Figure 55: Quota Templates in File Server Resource Manager

File Screening Management

This function allows you to prevent users from saving certain types of files to disk. It also comes with a set of pre-defined templates. For example, if you want to prevent users from filling up their hard drive space with illegal music or movie files, you can use the Block Audio and Video Files template.

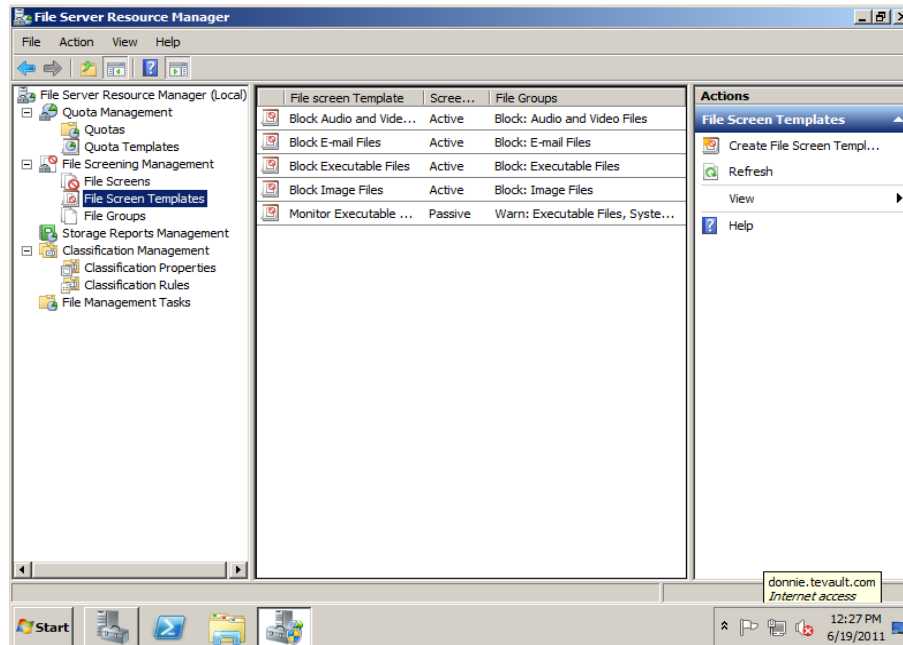


Figure 56: File Screen Templates in File Server Resource Manager

Storage Reports Management

The third major function of File Server Resource Manager is to prepare reports about file and disk management. You can either create reports on demand or have them automatically created on a regularly scheduled basis. Reports can be customized to your liking, with a variety of different criteria on which to report. In addition to information about quota usage, you can also generate reports that show what the largest files are, which files have been accessed the least, which files have been accessed the most, which files are owned by a particular user, and the like.

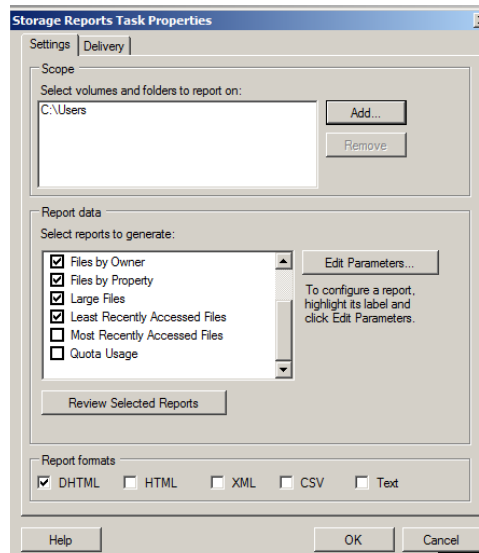


Figure 57: Creating a Storage Report in File Server Resource Manager

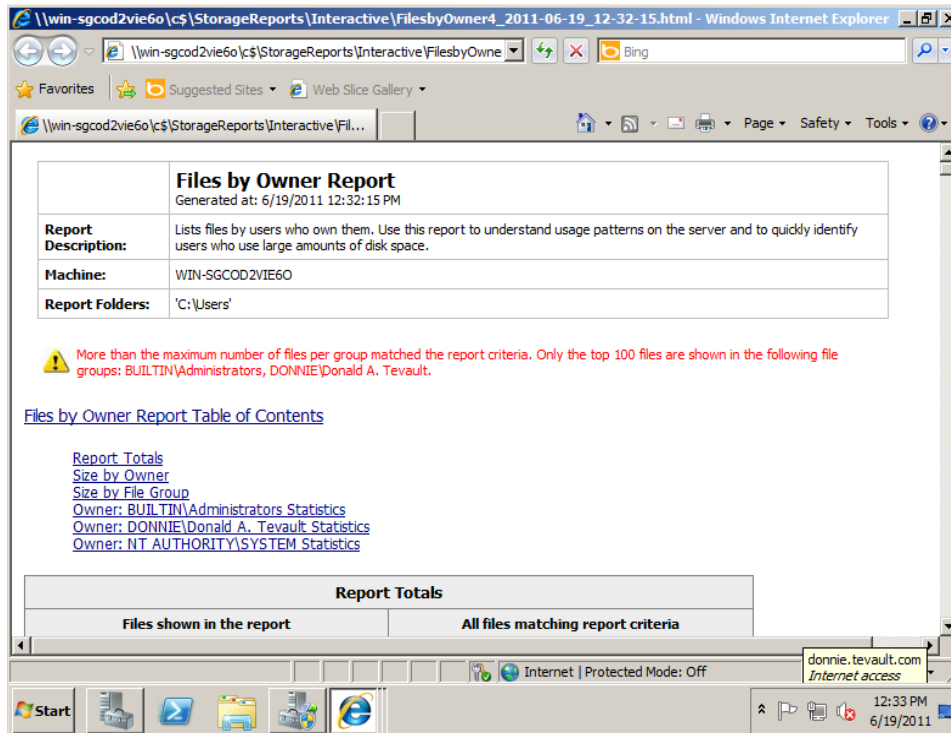


Figure 58: A Storage Report created by File Server Resource Manager

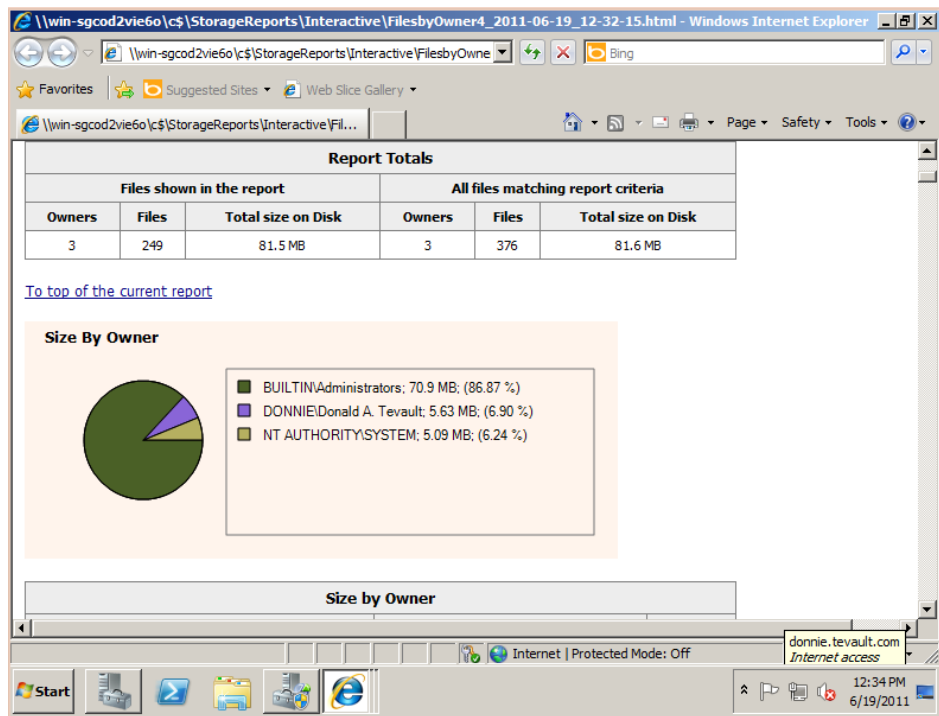


Figure 59: A Storage Report created by File Server Resource Manager

Security

When planning for setting up a corporate file server, it's important to keep security in mind. You want authorized users to have their required access to files and folders; however, you want to keep unauthorized people out. With Windows Server 2008, you can do this by way of share permissions, NTFS permissions, and the Encrypting File System.

Share Permissions

Share permissions determine how a folder gets shared across the network. Exactly what the Share permissions are depends upon whether you're looking at Windows Server 2008 R2 or the original Windows Server 2008.

For Windows Server 2008 R2, the Share permissions are:

- **Owner** – This person has full control over the folder.
- **Read** – Users can read the contents of the folder, but they can't write to it.
- **Read/Write** – Users can read and write to the folder.

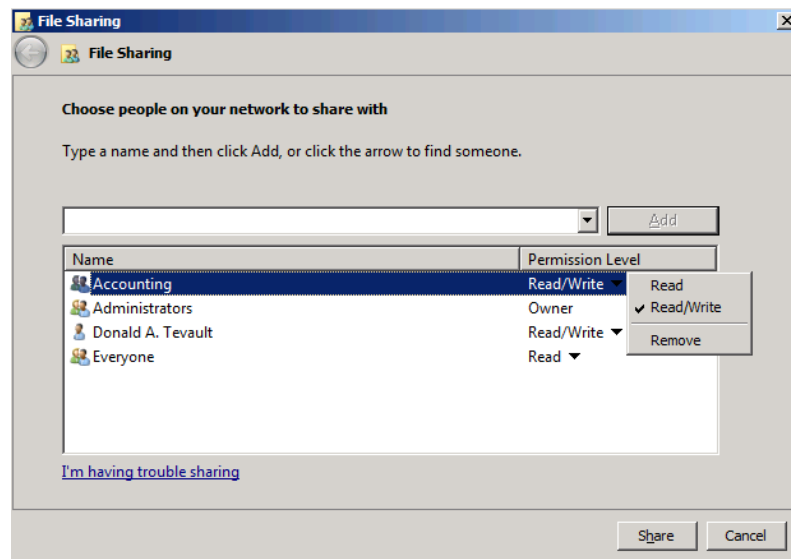


Figure 60: Share Permissions on Windows Server 2008 R2

For some strange reason, Microsoft had different names for these Share permissions in the original version of Windows Server 2008. In addition to Owner, there are:

- **Reader** – Users have the ability to read the share, but not write to it.
- **Contributor** – Users have both read and write permissions for the folder.
- **Co-Owner** – Has full control over the folder. This includes the ability to change permissions for the folder.

NOTE: Be familiar with both the Windows Server 2008 R2 and the original Windows Server 2008 Share permissions. We can't guarantee which version you'll see on the exam.

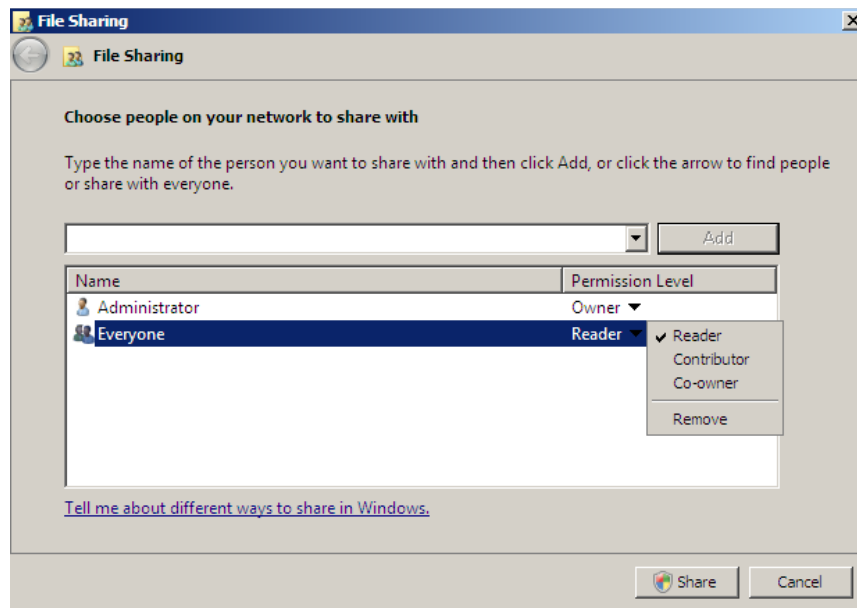


Figure 61: Share Permissions on Windows Server 2008 (Original)

To share a folder in a Windows Network, do as follows:

1. Make a new folder in a location of your choosing on your Windows Server.
2. Right-click the folder and choose **Share...**
3. On the file sharing permissions screen, enter the name of the users that you would like to give share permissions.
4. Click **Add**.
5. Set the share permissions level by clicking the drop-down menu.
6. Click **Share**.
7. When the share completes, click **Done**.

NOTE: Remember that Share permissions work at the network level.

NTFS File Permissions

Unlike Share permissions, NTFS permissions work at the file system level. The basic permissions settings are:

- **Full Control** – This setting allows users to perform any action at all on the file or folder.
- **Modify** – This setting allows users to read, edit, or delete files or folders.
- **Read & Execute** – This setting allows users to run an application.
- **List Folder Contents** – This allows users to browse the contents of a folder, but it doesn't necessarily allow them to open any files that are in the folder.

- **Read** – This allows users to browse the contents of a folder, and to open files that are in it. However, users won't be able to edit any files or run executable programs that are within the folder.
- **Write** – Users can create files within a folder, but they won't necessarily be able to read them once they're created. This is good for when multiple people have access to a folder, and you don't want users to see each others' files.

To change NTFS permissions settings, go to the Security tab of the Properties page, and click on the Edit button.

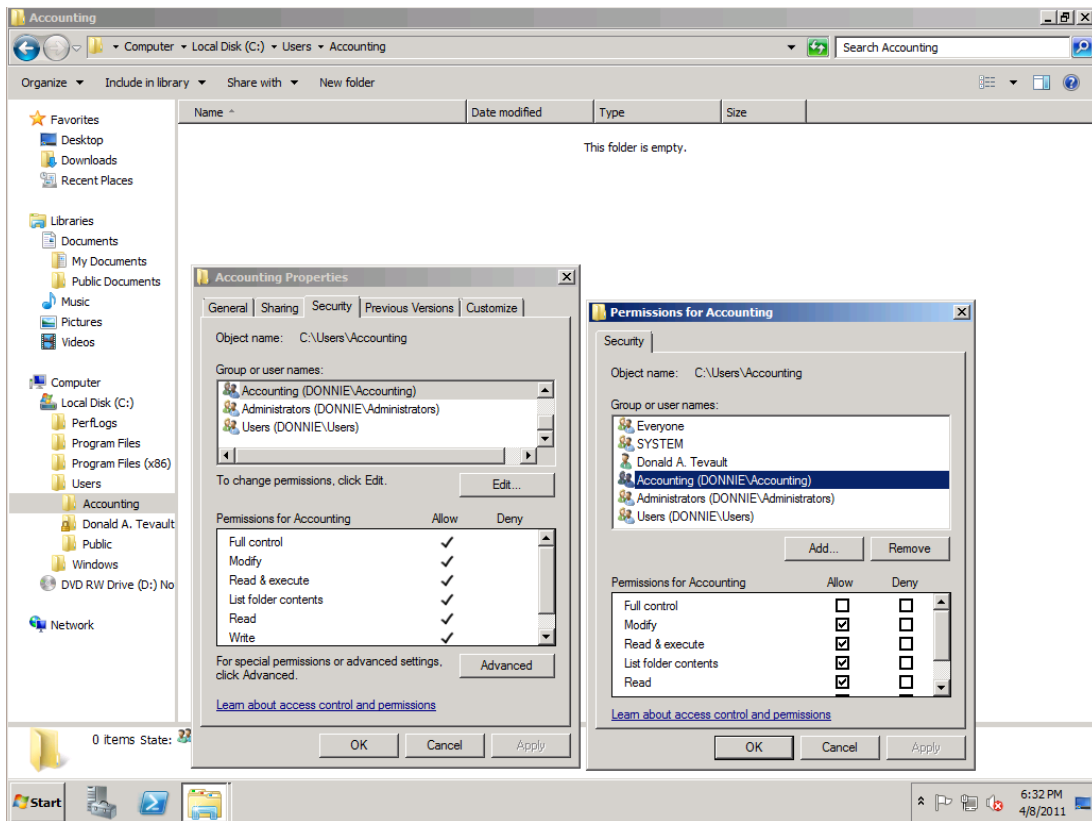


Figure 62: NTFS Permissions

Near the bottom of the Security tab, you'll see the Advanced button. By clicking on that, you'll be able to access a set of much more granular permissions settings. You can see a description of these settings in the chart below.

Special Permissions	Full Control	Modify	Read & Execute	List Folder Contents (folders only)	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					
Synchronize	x	x	x	x	x	x

Figure 63: Special Permissions

Encrypting File System (EFS)

A slight shortcoming with the NTFS permissions is that they only work as long as the computer is running. A malicious user could bypass all NTFS settings simply by booting the computer from a live CD and mounting the hard drive, or by taking the hard drive and mounting it in another computer. However, by using the Encrypting File System to encrypt files and folders, this shortcoming can be remediated. Anyone who tried to use one of these methods to bypass the NTFS permissions wouldn't be able to read the encrypted files.

EFS can be configured either for individual computers or for groups of computers via Group Policy. Data Recovery Agents, or DRAs, can be used to recover encrypted files in case the encryption keys are lost.

BitLocker

BitLocker is another way to encrypt data on a computer's hard drive. Where EFS is used to encrypt only individual files at the user's discretion, BitLocker will encrypt the entire hard drive. It also does integrity checking of the operating system's boot components, and can work together with a Trusted Platform Module to prevent unauthorized users from booting the computer.

While BitLocker is useful in the corporate server room, it's even more useful for branch offices where physical security for servers is usually not quite as strong. Even if a BitLocker-encrypted drive is inserted into another computer, a would-be data thief wouldn't be able to obtain any data.

As is also true with EFS, BitLocker supports the use of Data Recovery Agents. This allows administrators the ability to access drives that have been encrypted with BitLocker.

Also, BitLocker can be used to encrypt portable devices, such as USB Flash drives. Microsoft calls this "BitLocker-to-Go."

Data Recovery Agents (DRA)

Data Recovery Agents are used to recover lost encryption keys for either BitLocker or EFS. DRAs require that a properly configured Certificate Authority be set up, and that the Group Policy be properly configured. This can be done under the “Public Key Policies” section of the Group Policy Management Editor.

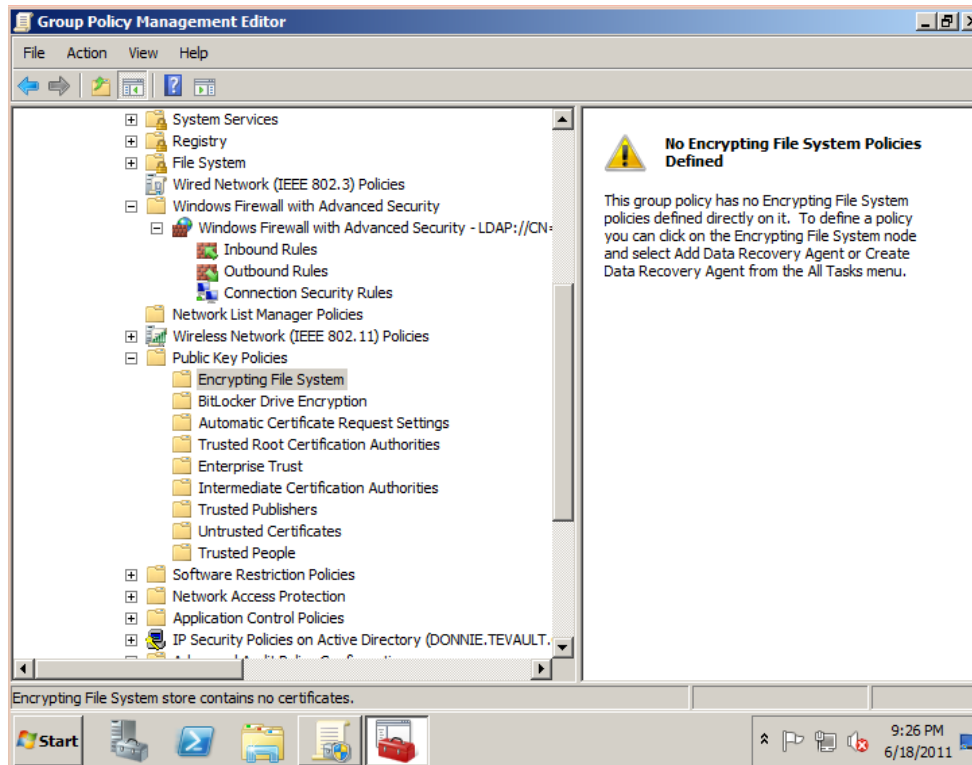


Figure 64: Public Key Policies

Publishing a File Share

File Shares can be published in Active Directory Domain Services. Doing so allows them to appear as a public resource within the Active Directory Domain Services infrastructure. To do so, you can do the following:

1. Open **Active Directory Users and Computers**.
2. Expand the Domain you would like to have the Shared Folder in.
3. Select the OU that should contain the shared folder.
4. Right-click in the right pane and select **New Shared Folder**.
5. Give the Shared Folder a public name, such as “PublicShare.”
6. Type the UNC path of your server (\\server\share).
7. The folder is now published in Active Directory Domain Services.

Once a shared folder is published, users can look for the shared folder both at the network share itself and within Active Directory Domain Services by using Active Directory resource management.

BranchCache

BranchCache is a brand-new feature that was introduced in Windows Server 2008 R2. It's a way to conserve Wide Area Network bandwidth by caching copies of documents at a corporation's branch offices. After documents on the central server have been cached at the outlying offices, clients no longer have to consume WAN bandwidth in order to access them. Note, though, that for BranchCache to work, the servers need to run Windows Server 2008 R2, and the clients need to run Windows 7.

There are two ways to set up BranchCache. With Distributed Cache Mode, there's no caching server at the branch office. Rather, each Windows 7 client will cache documents on its own local hard drive. Then the clients will send multicast messages to other clients at the branch office via Web Services Dynamic Discovery. This way, the clients can discover if documents have been cached on other local clients. In effect, Distributed Cache Mode sets up a peer-to-peer sharing system at the branch office.

With Hosted Cached Mode, documents from the central server are cached on a server at the branch office. (Note that this server must be running Windows Server 2008 R2, and must have the BranchCache feature enabled.) Then the Windows 7 clients at the branch office will access the local copies.

Either way, you would enable BranchCache on the central server by adding the File Services role, and then adding the BranchCache feature. The next steps would depend upon which caching mode you wish to use.

For Distributed caching, you would next need to configure Group Policy to "Allow Hash Publication Only for Shared Folders on Which BranchCache is Enabled." Then you would need to specify a "HashStorageLimitPercent" registry value. After that, you can enable BranchCache support for the files and folders you wish to share.

For hosted caching, you would need to open a command-prompt in administrator mode, and run the command, "netsh BranchCache set service mode=HOSTEDSERVER." Then install an SSL authentication certificate with the Fully Qualified Domain Name of the hosted cache server. Finally, obtain the hash value from the certificate that you just installed, and use it to execute the command, "netsh HTTP ADD SSLCERT IPPORT =0.0.0.0:443 certhash='whatever_your_hash_is' APPID={d673f5ee-a714-454d-8de2-492e4c1bd8f8}:"

The Windows 7 clients can be configured either by running a "netsh BranchCache" command or by editing Group Policy.

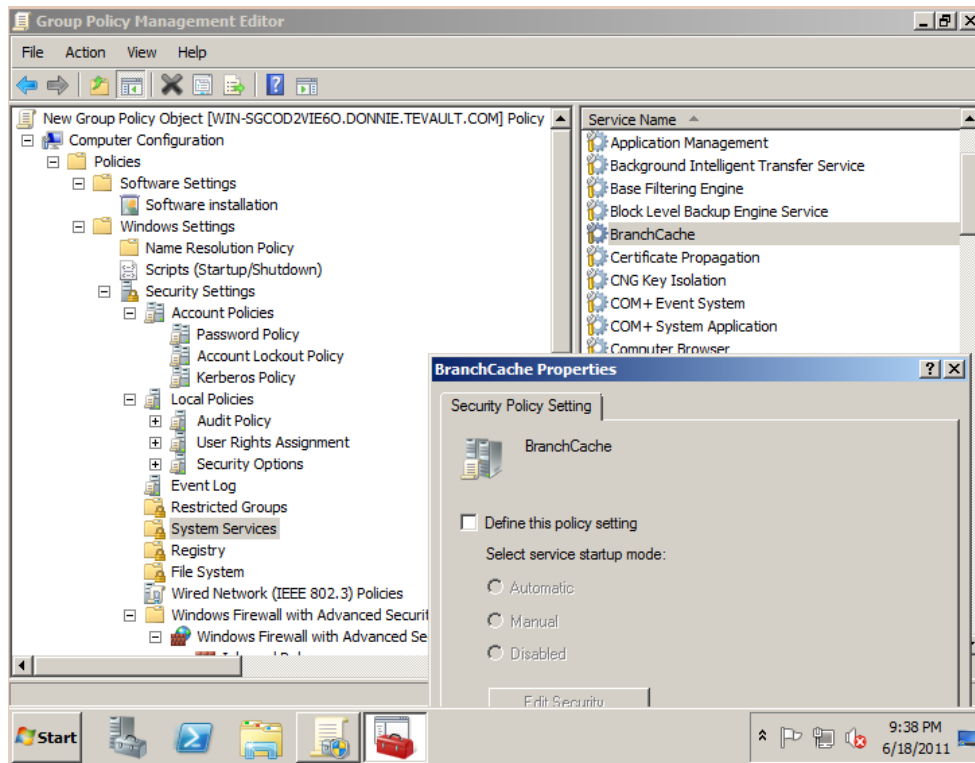


Figure 65: BranchCache Policy

Configuring Distributed File System (DFS)

DFS uses multiple servers to manage a file sharing network. It takes advantage of the resources of multiple servers that may either be in one location, or scattered across a Wide Area Network. Shared folders are automatically replicated to all servers in the DFS namespace. That way, if one DFS server goes down, the shared folder will still be available to users. Using a DFS namespace also allows users to access files without regard to which physical server holds the sought-after files.

When files change, DFS can replicate only the changed portions to other servers in the namespace. This can make for a considerable savings in network bandwidth.

With Windows Server 2008 R2, DFS can also be set up on a fail-over cluster, to ensure that data will always be available when needed. Windows Server 2008 R2 also allows you to set up shares with read-only permissions. This can be handy for when you need to make a DFS share available to a branch office location, but you want to ensure that nobody modifies the files on that share.

To set up DFS, do the following:

1. Open **Server Manager**.
2. Select **Roles**.
3. Scroll down to **File Services**.
4. Choose **Add Role Services**.
5. On the screen below, select **Distributed File System, DFS Namespaces**, and **DFS Replication**.

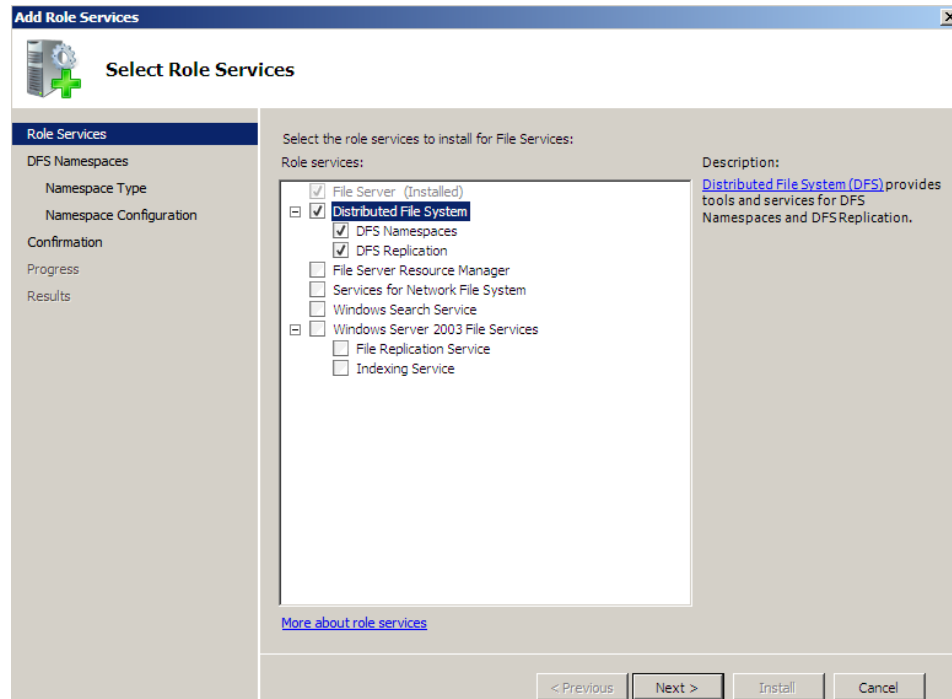


Figure 66: Adding DFS Role Services

6. Click **Next**.
7. On the next screen, you can choose your namespace. You can either create a new namespace now or do it later by using the DFS snap-in.
8. Choose to create one now and name it something appropriate, like DFSshare.
9. Click **Next**.
10. On the next screen, you can choose from a Domain-Based or Namespace-Based installation.
11. For our example, choose Domain-Based and make sure the 2008 mode checkbox is checked.
12. On the next screen, partially shown below, you can choose the name of your shared folder target. Here we have chosen DFS share, but you could alternatively add extra folders using the **Add** button.

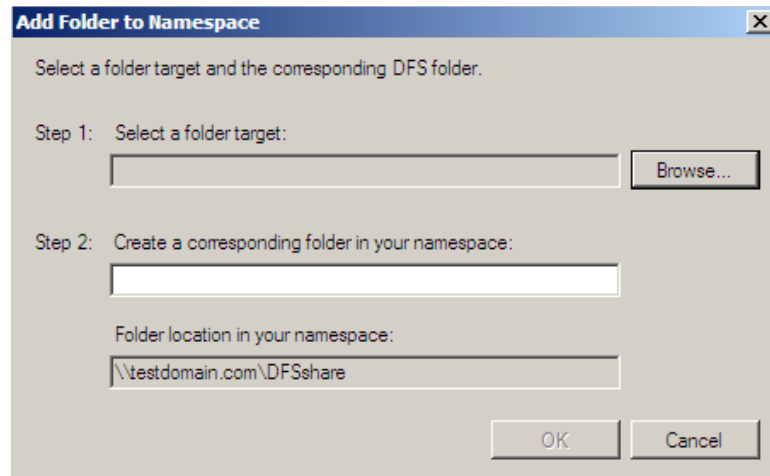


Figure 67: Adding a Folder to a Namespace

13. You can add a folder into the namespace that will be accessible whenever a user accesses the namespace.
14. Click **Next**.
15. Click **Install**.

You can access your new DFS share by going to **Administrative Tools > DFS Management**. Here you can create new namespaces and replication groups, or you can modify the DFS share. We will now show how to create a replication group. Replication groups are designed to create a second namespace that balances the content between two servers.

To create a replication group, do the following:

1. From **DFS Management**, select **New Replication Group**.
2. From the first menu, select either **Multi-Purpose Replication Group** or **Replication Group for Data Collection**. In this example, we'll create a multipurpose replication group. Click **Next**.
3. Name the replication group. We're naming ours "**Test**."
4. Make sure that the Domain is set correctly and click **Next**.
5. Select a member server and another server that will replicate, then click **Next**.
6. Under the topology, you can choose **Hub and Spoke**, **Full Mesh**, or **No Topology**. A brief description of each is available in the console, as shown in the screenshot below:

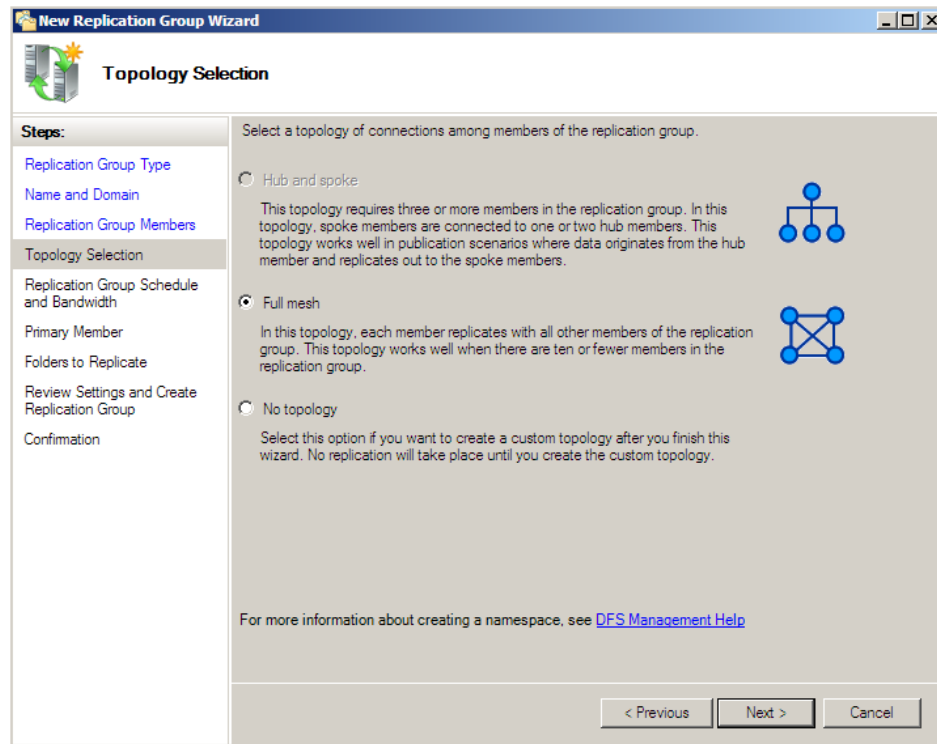


Figure 68: Choosing a Topology for DFS Replication

7. We'll choose **Full Mesh** in our example, where each server replicates with each other. Click **Next**.
8. In the replication schedule, you can choose the amount of bandwidth used and on which days. Normally, you would just pick **Full**. However, you can choose from varying bit rates, from 16Kbps to 256Mbps. Click **Next**.
9. At the next screen, choose the primary member server. This will be the server which is the "main" server from which the files are replicated. Click **Next**.
10. On the next page, choose a logical folder to replicate. For example, a shared folder on your hard drive called "Share." Enter the full drive path. Click **Next**.
11. Click **Next**.
12. Click **Create**.
13. A series of tasks will flash by in the confirmation screen. When it is finished, you will need to click **Close**. The DFS replication is now set up.

Configuring Volume Shadow Copy Services (VSS)

Volume Shadow Copy Services allow you to back up files that are open and in use. It's easy to set up and even more easily used. To set up VSS, do the following:

- From the computer Explorer window, right-click a system volume and choose **Configure Shadow Copies**.
- Select your hard drive and choose **Enable**.
- Click **Settings** and choose "no limit" to have a full shadow copy.
- Click **Schedule** to set a schedule for the replication of shadow copies. By default, the shadow is copied twice per day. You can set it in the menu you see below:

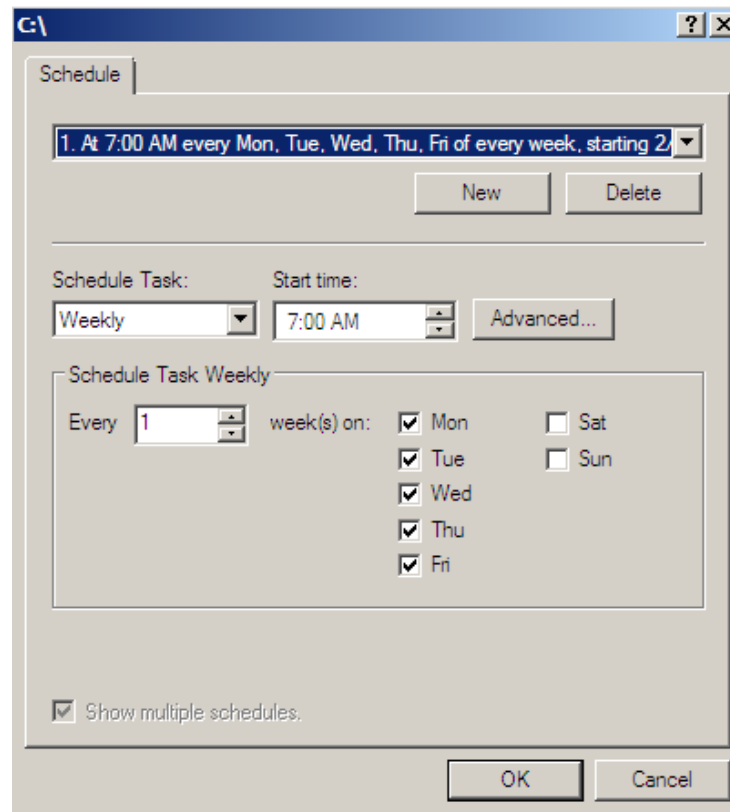


Figure 69: Scheduling the Volume Shadow Service

- As the warning message says, you should avoid making more than two shadow copies per hour.
- Lastly, on the **Details** menu in the shadow copy menu, you can see the amount of space used by your shadow copies.

You can also create shadow copies from the command-line by using the VSSAdmin utility.

Configuring Backup and Restore

Unlike its predecessors, Windows Server 2008 has backup software built right in. It's still not perfect, since its capabilities to perform backups over the network are rather limited, but at least it's better than what Windows Server had before.

To install Windows Server backup, follow these steps:

1. Open **Server Manager**.
2. Select **Features**.
3. Select **Add a Feature**.
4. Choose "**Windows Server Backup Features**."
5. Click **Next**.
6. Click **Install**.

Once installed, Windows Server backup can be launched from the Server Manager. You can either choose to run a one-time backup, or you can schedule regular automatic backups. With Windows Server 2008 R2, you can also schedule a backup to another location on the network.

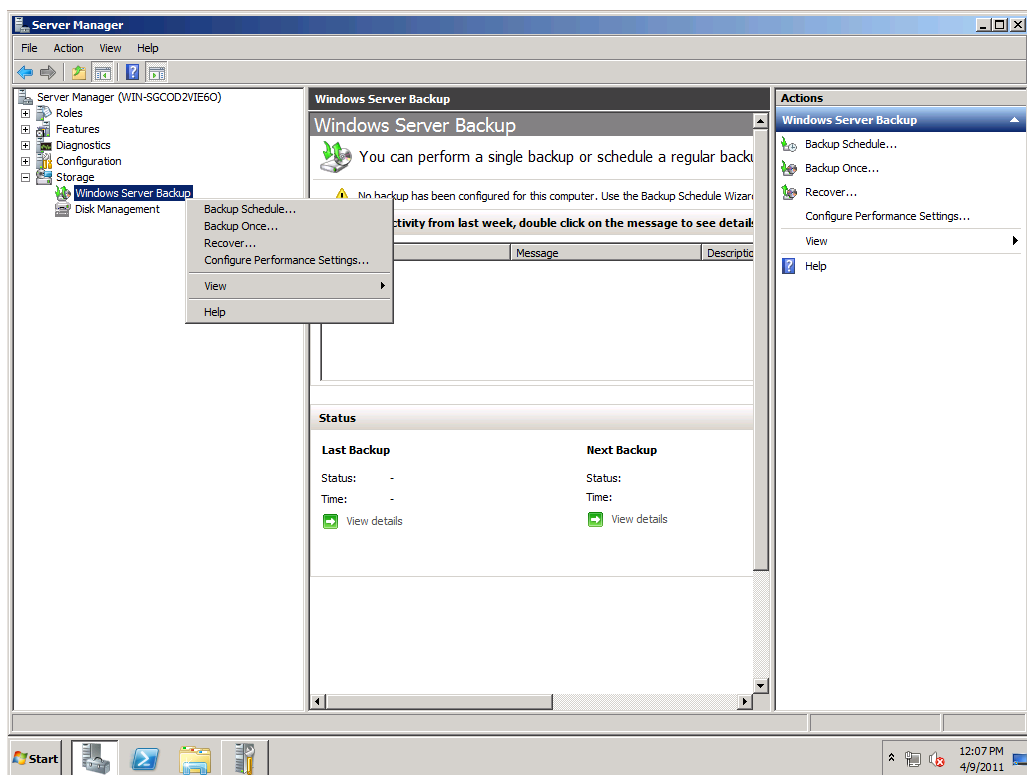


Figure 70: Windows Server Backup

Managing Disk Quotas

Even if you haven't installed the File Server Resource Manager, you can still set disk usage quotas to prevent users from using too much drive space. When you use this method, quotas will be set on a per volume basis.

To setup a quota, do as follows:

- Navigate to a disk volume, right click, and select **Properties**.
- From the menu you see below, check **Enable quota management** and the **Limit disk space** button. There, you can enter quota limits. We've set it to 100 Megabytes.

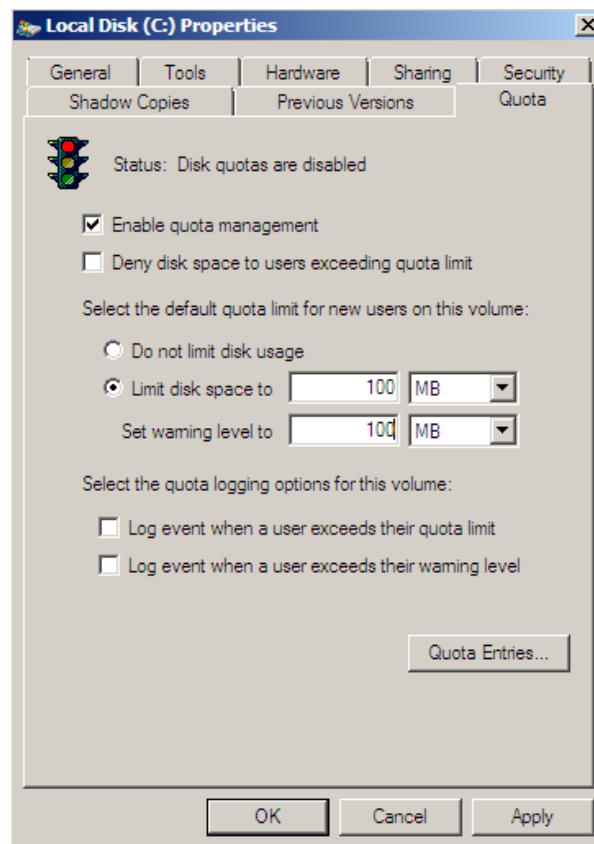


Figure 71: Setting Up a Disk Quota

- Placing this 100MB quota limit sets a default quota for all users in the domain.
- However, should you wish to specify a specific quota, you can click **Quota Entries...**
- This will bring up the quota entries screen you see below.




Figure 72: Specifying Quota Entries

- Selecting **Quota > New Quota Entry...** will bring up a specific quota entry, as shown below.




Figure 73: Adding a New Quota Entry

- For administrator accounts, you can override the default setting. You can either allow the administrator account to be free of a quota, or you can set the administrator account with a higher quota than other users.

Configuring and Monitoring Print Services

Configuring Printers with the Graphical Interface

One area where Windows Server truly excels is the management of printers and print resources. Using Windows Server's print management utilities, you can administer from dozens to hundreds of printers through printer shares, connections, and other methods. In this section, we will explore some of these methods.

First, to publish a printer share:

1. Install a printer through the drivers associated with that printer.
2. Navigate to **Control Panel**.
3. Right-click the printer and choose **Sharing**.
4. Set the sharing properties you would like to set for your printer, based on activity directory groups.

To install File Printer Services, do as follows:

1. Open up **Server Manager**.
2. Select **Roles** and choose **Add Roles**.
3. Choose **Print Services** and click **Next**.
4. Click **Next** again, then choose **Print Server, LDP Service** (Line Printer Daemon), and **Internet Printing**. This sets up a website that allows users to manage print jobs.
5. Click **Next**.
6. Click **Next** again, and then click **Install**.

To add a printer driver to the print server:

1. Open **Print Management** by navigating to **Administrative Tools > Print Management**.
2. Click **Print Servers** and then click **Printers**.
3. Right-click the printer you want to add additional drivers to, then click **Manage Sharing**.
4. Click **Additional Drivers**.
5. Select the process architecture, such as x86 or 64 bit.
6. You can choose the location of the print drivers in the final menu. **Note:** You can also update and change printer drivers through this menu.

To publish a printer in Active Directory Domain Services:

1. Open **Active Directory Users and Computer** by navigating to **Administrative Tools > Active Directory Users and Computers**.
2. Choose your Computer's Organizational Unit (or an OU that you don't mind having a printer resource in).
3. Right-click the action pane and choose **New > Printer**.
4. Enter the network path of the printer.
5. Click **OK**.

Configuring Printers on Windows Server 2008 Core

For a server that's running Windows Server 2008 Core, you can use the "Pubprn.vbs" script from the command-prompt to publish printers to Active Directory. To import printers from other servers, you can use the "printbrm" utility.

Printer Pools

By adding groups of identical printers to a printer pool, you can allow print jobs to go to whichever printer isn't busy with another job. It also allows for redundancy, in case one printer has operational problems.

Working with Unix Servers

By installing the LPD Service and configuring the printers to use Line Printer Remote Printing, you can allow any Unix-based servers on your network to use the printers that are connected to your Windows Server machine.

Monitoring

You can use the **Reliability and Performance Monitor** to monitor printer performance:

1. Open **Performance Monitor** by navigating to **Administrative Tools > Performance Monitor**.
2. Under **Monitoring Tools** you can define performance and reliability views.
3. You can define data collector sets, such as Spooler Default Session.
4. You can create a custom filter that will send a status email if the printer runs out of paper or suffers a paper jam.

File Server Resource Manager

You can also create reports on disk system usage with the File Server Resource Manager. We've already covered this topic in the "Configuring a File Server" chapter.

Domain 5: Monitoring and Managing a Network Infrastructure

So far, we have covered the process of setting up a network infrastructure and ensuring that the infrastructure is well-designed and can support our current number of users and any future employees we may acquire. Next, we come to one of the most important parts of any network design: **maintenance**. In this section, we'll address how to monitor the behavior of our network and determine if maintenance needs to be done.

Configuring Windows Server Update Services (WSUS)

Windows Server Update Services is a software update method utilized by Microsoft to automatically deploy updates from a server to its child servers and workstations. This ensures easy, centralized administration and management of any Microsoft update to every system on your enterprise network.

Properly implementing this feature requires a mix of group policy, targeting specific clients, testing the implementation, and preparing for when networks may disconnect from time to time. As a first step, we need to install WSUS.

To install WSUS, do the following:

- Confirm that your Server contains at least 1 GB of RAM per 500 clients.
- Prior to Windows Server 2008 SP 2, you would have to download WSUS with the latest services pack from TechNet, and then perform a manual installation. With either Windows 2008 SP2 or Windows Server 2008 R2, all you need to do is to open the Server Manager page, and add WSUS as a server role. This will automatically download and install WSUS for you.
- Add the **Application Server** and **Web Server (IIS)** roles to your server. You can do this by doing the following:
 1. Start **Server Manager**.
 2. Select **Roles** and then **Add Roles**.
 3. Choose **Application Server** and **Web Server (IIS)** and agree to any additional prerequisites.
 4. Ensure that **IIS6 Compatibility** is checked when you install IIS.
- Either double-click the download or use the **Server Manager** to install WSUS 3.0. To use the server manager, do the following:
 1. Start **Server Manager**.
 2. Click **Add Roles**.
 3. Choose **Windows Server Update Services**.
 4. Confirm the selections and click **Install**.

Once WSUS has installed, you will need to configure it by doing the following:

- Accept the license agreement.
- Agree to install any additional or missing components.
- Select an **Update Source** location, as you see in the following screen:




Figure 74: Choosing a WSUS Update Source

- Choosing the default location is usually sufficient. If you like, you can specify a folder somewhere on your server or an NFS/SMB location.
- Next, choose where to keep an internal database, as you see in the screen below. It's highly recommended that you pick the same directory that you stored the physical updates in. (**Note:** This database is not an SQL database. To select an SQL database, you will need to select either the existing database or another existing database and provide the **machinename\instancename**.)

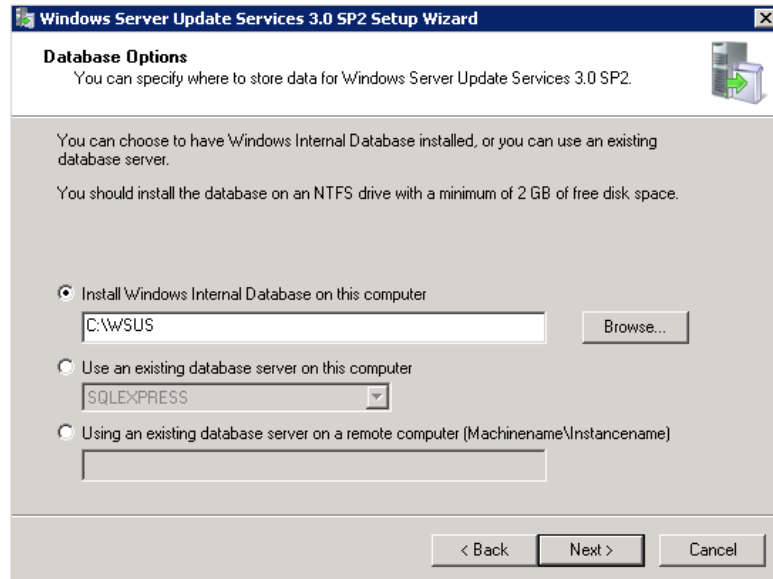


Figure 75: Choosing an Internal Database for WSUS

- At the **Web Site Selection** screen, you will want to choose to either create a new website or choose an existing website. As a best practice, you should create a new website by selecting the new web site radio button. This is a best practice because it exercises one of the fundamental aspects of good administration: separation. It also practices another great habit: dedication. It's always good to have a dedicated, separate server for a specific task.

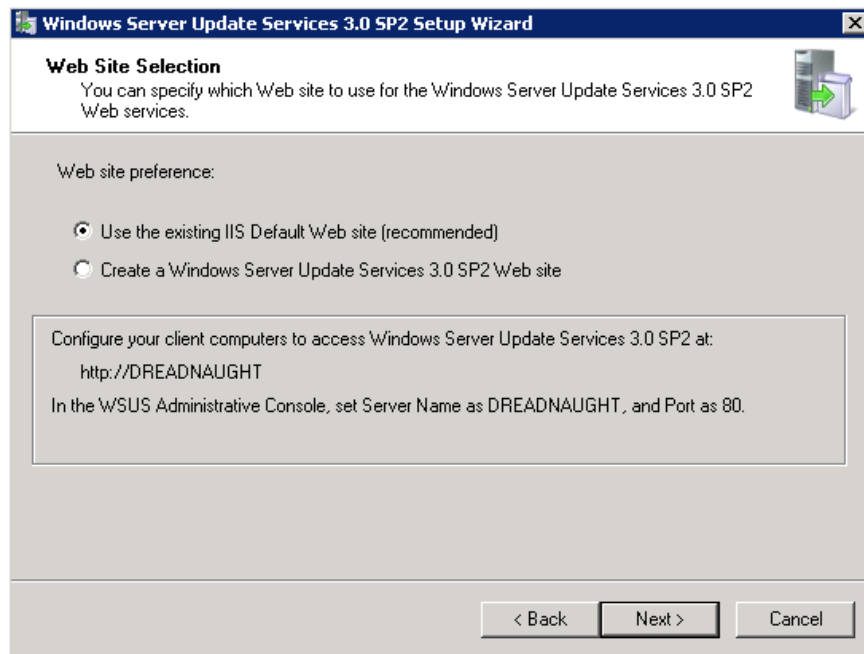


Figure 76: Choosing a Web Site for WSUS

- Click **Next** at the following two screens and wait for the components to install.
- Click **Finish**.

Now you will need to configure WSUS to apply server updates to your infrastructure. To do so, follow these instructions:

1. Click **Next** at the opening Wizard.
2. The next screen will ask you to participate in Microsoft's improvement program, you can choose to participate or not. Click **Next**.
3. At the next screen, shown below, you can choose to **Synchronize from Microsoft Update** or to **Synchronize from another Windows Server Update Server**.

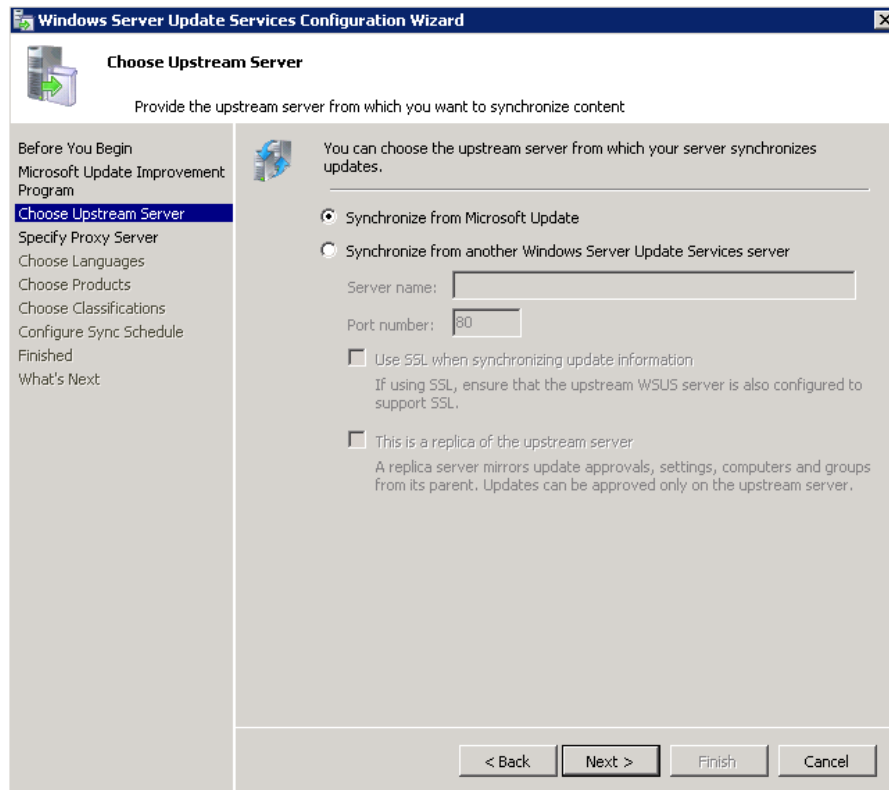


Figure 77: Choosing the Upstream WSUS Server

This portion of the WSUS service is designed to give you the ability to tier updates. Instead of having each computer on the network downloading updates from Microsoft, you would just have one server do it. Then you could have that server send updates to other downstream servers and clients in your organization, choosing which updates you would like to install on them. For the purposes of this example, we will use Microsoft Update, thus creating the upstream server for our organization. On subsequent servers, you can fill in the server name and port at this step.

Take note of the choices provided under the server name and port number fields:

Using **Secure Sockets Layer (SSL)**, you can force client computers and downstream servers to authenticate the WSUS server before taking updates. Additionally, SSL will encrypt any metadata passed between client and server computers.

You may also mark the server as a **replica**, which mirrors updates, settings, and groups from the parent update server, providing high availability.

1. If you have a **proxy server**, you will need to enter it on the next screen. If not, just click **Next**.
2. At the next screen, you'll choose the types of updates available and the languages you will make your updates available in.
3. Click **Start Connecting**. The portion in the updates area may take some time.
4. Once the process completes, click **Next**.

5. In the language options menu, shown below, you can choose the languages you would like to receive updates in. If you're an organization with employees who natively speak another language, you will want to choose languages for these individuals here. Once you have made your selections, click next.

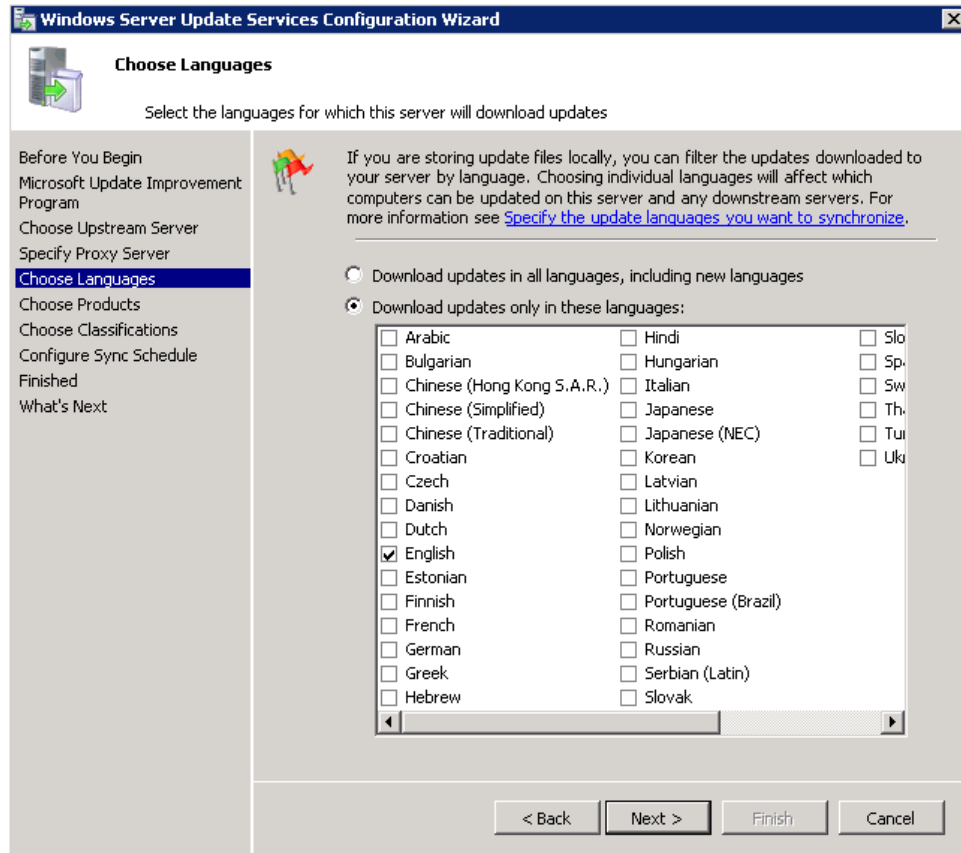


Figure 78: Choosing Update Languages

6. The next screen allows you to select the products for which you would like to receive updates. This includes non-operating system products, like Office and SharePoint. According to Microsoft, WSUS can update everything in your organization easily. Make your selections and click **Next**.

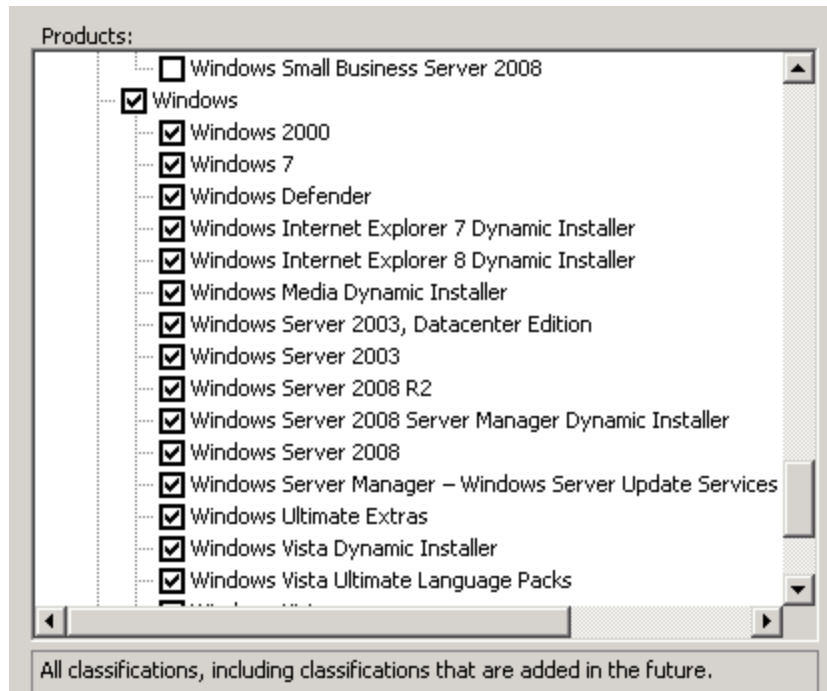


Figure 79: Choosing Updates for Specific Products

- The next screen, shown below, allows you to choose classes of updates. Best practices dictate that you enable critical updates, definition updates, and security updates, at least. Make your selections and click **Next**.

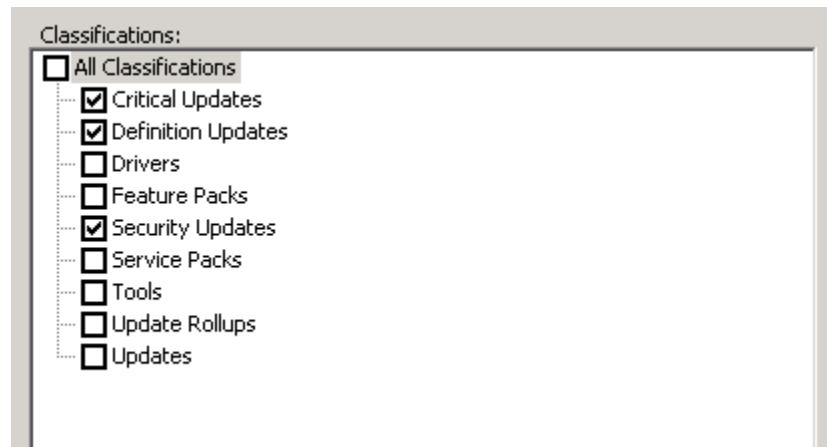


Figure 80: Selecting Update Classifications

- At the next screen, you can choose to manually or automatically sync your updates. Manually syncing will allow you to choose when your servers check in for updates. The servers will download updates depending on your approval. Automatically syncing will cause them to update on a regularly scheduled basis. In this example, we've chosen automatic updates at 4:17:08AM, once per day.

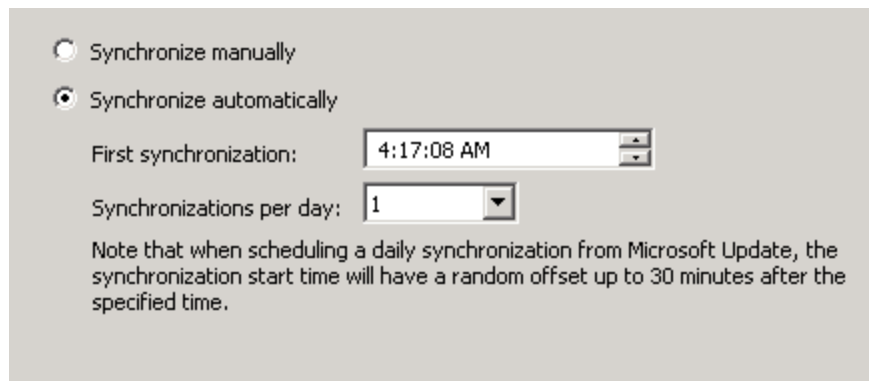


Figure 81: Configuring WSUS Synchronizations with Microsoft

9. Click **Next**.
10. At the next screen, you will begin your initial synchronization. Run this now. You can either click **Finish** and go to the next screen or click **Next**.
11. If you click **Next**, you will see more advanced options, including the following:
 12. **Using SSL with WSUS** – Configuring your updates to be sent from a secure location. Remember that you will need two ports: one for encrypting metadata and one for clear HTTP transmission.
 13. **Create Computer Groups** – Create security groups to deploy WSUS to.
 14. **Assign Computers to Groups using Group Policy** – Configure group policy for WSUS groups.
 15. **Configure Auto-Approval Rules** – Configure rules to auto-approve update pushes for your clients.
16. In the next section, we will walk through some of these submenus that pertain to the exam.

Creating a Computer Group

For various reasons, you may want to create several different computer groups for clients that you want to update from WSUS. You may, for example, create separate groups for all the different models of computers. Or you may just want to ensure that clients don't get updated all at the same time.

Selecting the Create a Computer Group button will take you to the Microsoft Help Menu file, which will take you through the steps of creating a computer group:

1. From the updates services tree, expand your server and then expand computers.
2. Under **All Computers**, choose the group where you would like to create a new computer group.
3. Right-click your node and select **Add Computer Group**.
4. Type the name of your group and hit **OK**.

Assigning Clients through Group Policy

After you create the groups, you'll need to manually add the appropriate computers to them. Then create the Group Policies by performing the following steps:

1. Open **Group Policy Management**.
2. Right-click the OU containing the computers in your organization and create a GPO. Edit the GPO.
3. Expand **Computer Configuration > Administrative Templates > Windows Components**, and select **Windows Update**.
4. Double-click **Enable Client-Side Targeting**.
5. Click **Enabled**.
6. Type the name of the computer group in the "Target Group name for this computer" box.
7. Click **OK**.

Here's a list of the Group Policy settings that pertain to WSUS:

- Specify Intranet Microsoft Update Service Location
- Configure Automatic Updates
- Automatic Updates Detection Frequency
- Allow Non-Administrators to Receive Update Notifications
- Allow Automatic Updates Immediate Installation
- Turn On Recommended Updates Via Automatic Updates
- No Auto-Restart for Scheduled Automatic Updates Installations
- Re-Prompt for Restart with Scheduled Installations
- Delay Restart for Scheduled Installations
- Reschedule Automatic Updates Scheduled Installations
- Enable Client-Side Targeting
- Enable Windows Update Power Management to Automatically Wake Up the System to Install Scheduled Updates
- Allow Signed Updates from an Intranet Microsoft Update Service Location
- Do Not Display "Install Updates and Shut Down" Option in Shut Down Windows Dialog Box
- Do Not Adjust Default Option to "Install Updates and Shut Down" in Shut Down Windows Dialog Box
- Remove Access to Use All Windows Update Features

Auto-Approval Rules

Auto-approval rules allow you to choose certain types of updates to be automatically installed without administrative approval. This isn't recommended practice for most updates. However, you may want to do it for critical updates or security updates. If you'd like to configure your server to auto-approve updates for your client workstations, follow these steps:

1. In the **WSUS Update Services** tree, expand the section where you'd like to approve updates and click **Options**.
2. Click **Automatic Approvals**.
3. Select **Default Automatic** approval and then click **Edit**.
4. Click **New Rule** to create a new rule.
5. Edit the **Properties** of the rule and select where you would like to approve updates.
6. Click **OK**.

Autonomous vs. Replica WSUS Servers

If you have more than one WSUS server on your LAN, you can cut Internet bandwidth usage by having only one WSUS server obtain its updates from the Microsoft Update site. The other WSUS server can then get its updates from the first server, which is known as the "upstream" server.

A "replica" WSUS server is a "downstream" server that inherits its approval information from the "upstream" server. This may be desirable when all departments in the enterprise have the same requirements for their updates.

If a certain department has different update requirements, then you can set up its downstream WSUS server to run in autonomous mode. That way, it won't inherit approval information from the upstream WSUS server, leaving the department administrators free to approve their own updates.

Capturing Performance Data

An important part of any maintenance plan is the consistent gathering of network and system performance data for use in baselines and comparison studies. Through Windows Server 2008, administrators are given access to many tools that can be used to determine the overall health of systems, and whether or not they are performing at the ideal level.

Additionally, as problems begin to arise in your infrastructure, you will want to collect data sets that will allow you to analyze what is happening.

Performance Monitor

Performance Monitor, or "PerfMon" for short, is a performance analysis tool that is used to monitor the overall performance of your server. To launch PerfMon in Windows Server 2008, you can type **perfmon** in the Windows Start menu. This will launch the **Reliability and Performance Monitor**, shown below.

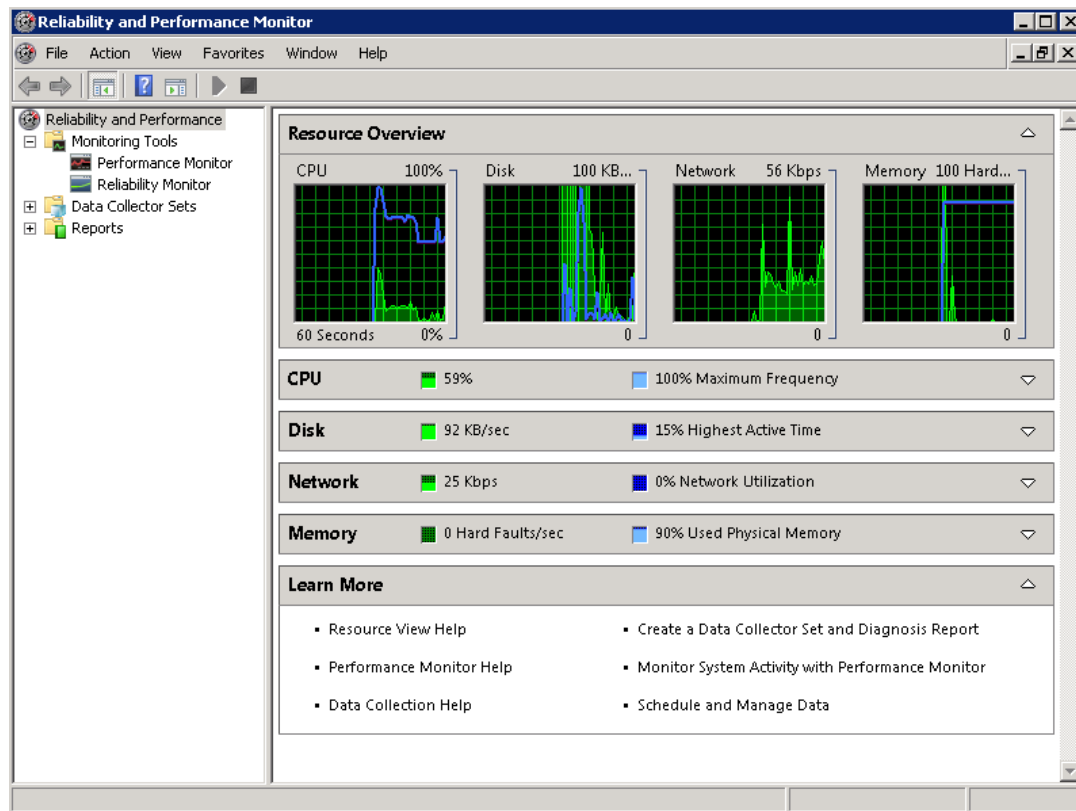


Figure 82: Windows Server 2008 Reliability and Performance Monitor

From the home screen, PerfMon creates an analysis of overall CPU, Disk, Memory, and Network use. Here we can get a quick overview of our system. In the screenshot above, we can see the following:

1. Our Current CPU load is 59%.
2. The Disk I/O is only 92Kbps.
3. The Network usage is only 25kbps.
4. There are no memory faults, but 90% of the memory is used.

The performance monitor section of the Reliability and Performance monitor can be accessed by selecting Performance Monitor from the main menu. In this section, you can create a number of fields to analyze, based on your custom inputs. To create a counter, you can click the plus sign and add many different fields to monitor.

For example, on a server in the infrastructure, we are interested in monitoring the performance of SQL. Therefore, in the “add counters” area, we can select the **MSSQL\$SQLEXPRESS:Databases** section and all the submenus beneath the plus sign. Then we can click **Add** and receive the added counters menu, as seen below.

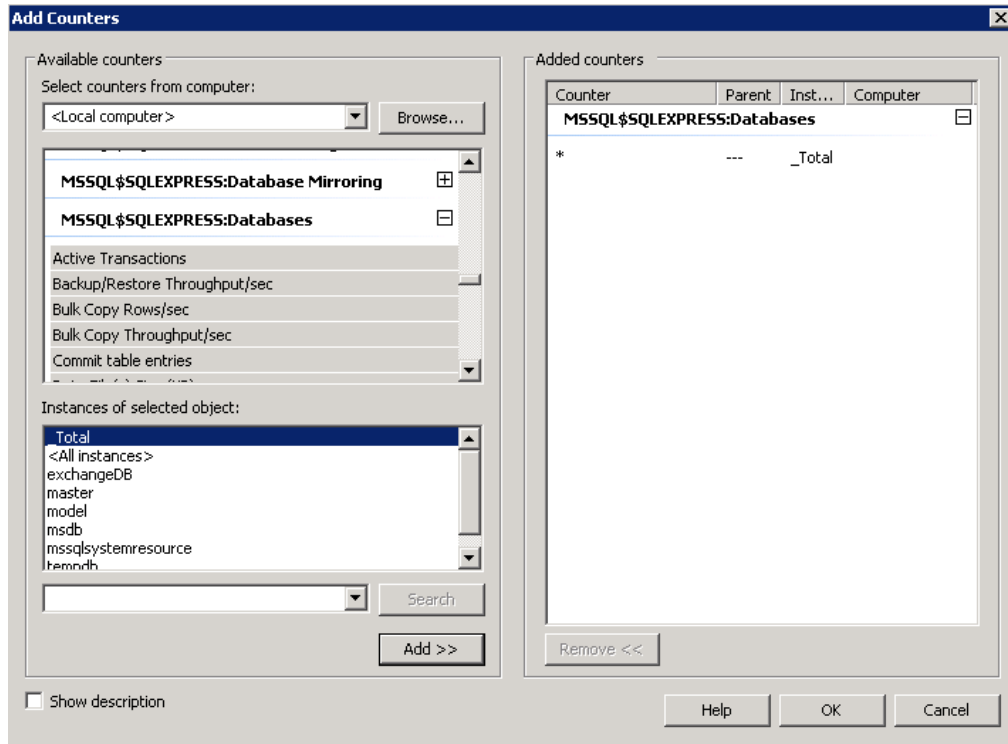


Figure 83: Adding a PerfMon Counter

Now when we hit **OK**, we will see the PerfMon counter that has been added, along with all the data fields associated with it. These will appear in the main screen, where we can watch their progress.

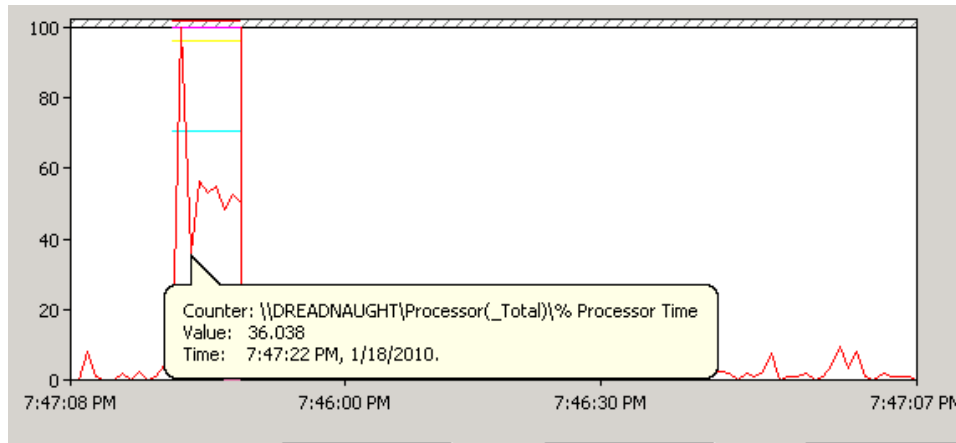


Figure 84: A Counter in Action

On the other hand, the reliability monitor will give you an overall view of your server’s health. Below, we have included a screenshot of the average performance of a server over time.

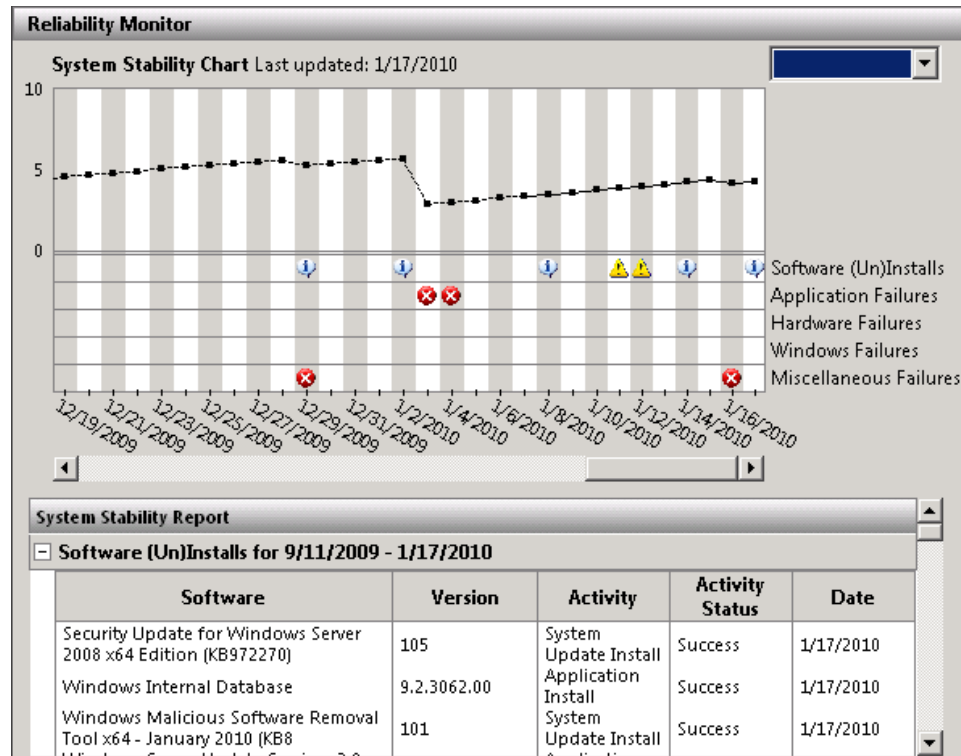


Figure 85: The System Stability Chart

The reliability monitor shows application failures, hardware failures, and other information pertaining to the server. From the drop-down menu, you can choose a date range or select all.

Data Collector Sets

Data Collector Sets allow administrators to create predefined rules that allow the collection of performance and reliability data. To create a data collector set, do the following:

1. Under **Data Collector Sets** in PerfMon, choose **User Defined**.
2. Right-click and select **New Data Collector Set**.
3. Name the data collector set. You can either create a collector set from a template or do it manually. It's much easier to choose a template, because most actions are predefined.
4. From the templates, you can choose:
 - ▶ Active Directory Diagnostics
 - ▶ Basic
 - ▶ System Diagnostics
 - ▶ System Performance
 - ▶ Health Registration Authority
 - ▶ LAN Diagnostics

5. For this example, we will choose system performance.
6. Click **Next**.
7. Keep the root directory default, unless you want to specify another directory, and click **Finish**.
8. This will monitor system events and collect them in the directory we defined in the previous step.
9. When the set appears in the user defined area, right-click it and select **Start**.
10. The set will begin to collect data and output it to the defined directory.

You can choose to view any of the pre-defined sets whenever you feel it is necessary. Each set allows you to pick various resources to view any time you wish. For example, if your LAN is acting strangely, choose LAN diagnostics, start it, and then monitor the results in the output directory.

Monitoring Event Logs

The Event Viewer is another powerful monitoring tool in Windows Server 2008. Through it, you can view a wide variety of events.

The event log is a very simple system to use. It can export data to XML files and “Event Files,” or EVT files. XML files are serialized files that contain data separated by tags. EVT files are specifically formatted log files that contain XML data, as well as event viewer data. EVT files also contain data associated with the Windows Event Viewer. With the Event Viewer, you can log Windows activities on five different levels:

1. Application
2. Security
3. Setup
4. System
5. Forwarded Events

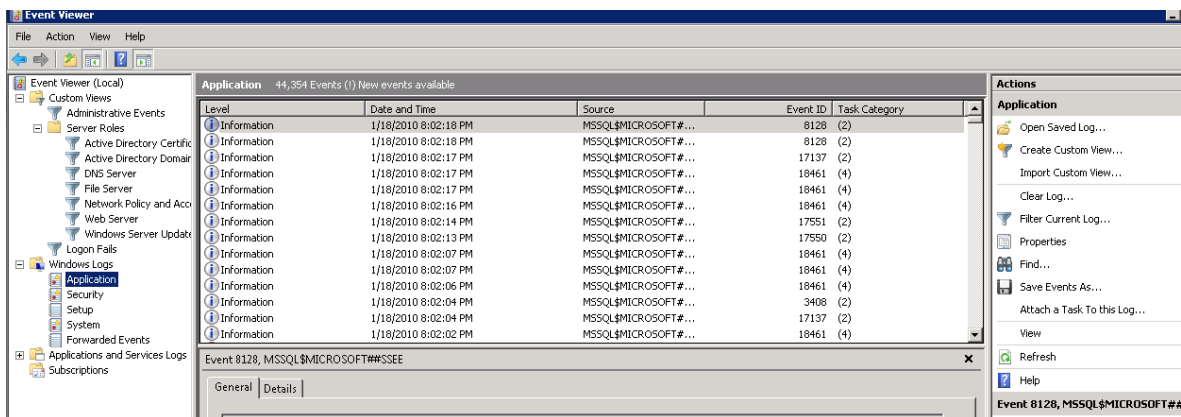


Figure 86: Windows Server 2008 Event Viewer

Regardless of which you choose, you can create custom views by clicking the **Create Custom View...** button and then configuring them from the very simple launch Wizard. Additionally, you can export these views by choosing **Save Events As...** and exporting them.

A relatively new feature of the Windows Event Viewer is the ability to consolidate log files on one central computer. This way, an administrator won't have to go all over the place to review computer logs. It also makes it easier for the administrator to keep an eye on things that could turn into problems, such as disk space on a fileserver. In addition to Windows Server 2008, Windows Vista, Windows 7, and Windows Server 2003 R2 can be configured to act as central log file repositories. Computers running Windows XP SP2 or SP3, Windows Server 2003 SP1 or SP2, Windows Server 2003 R2, Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 can be configured to send logs to a central log server.

To configure a Windows Vista, Windows 7, or either version of Windows Server 2008 to forward logs to another server, you'll need to open a command-prompt window with administrative privileges and enter the command:

```
winrm quickconfig
```

To configure a computer with Windows Vista, Windows 7, or either version of Windows Server 2008 to receive log files from other computers, open a command-prompt window with administrative privileges and enter the command:

```
wecutil qc
```

Alternatively, you can configure the receiving server by opening the Event Viewer and configuring event subscriptions. There are two types of event subscriptions.

- With a collector initiated subscription, the collecting computer will contact the source computers and request that they send it their new log events. With this type of subscription, the source computers must be members of the domain.
- With source computer initiated subscriptions, the source computers will contact the collecting computer. With this type of subscription, the source computers may or may not be members of the domain. If the source computers are not domain members, they must be configured with a valid certificate, and the CA that issued the certificate must be added to the collecting computer. Also, the source computers will have to be configured to contact the collecting computers. This can be done either via Group Policy or via local configuration.

To set up event subscriptions, you'll want to log into the collecting computer as a Domain Administrator.

Gathering Network Data

For the exam, network monitoring requires knowledge of Network Monitor configuration and practical use of the Simple Network Management Protocol (SNMP).

SNMP

The Simple Network Management Protocol, or SNMP, is a network protocol used by TCP/IP to monitor routers, bridges computers, wireless access points, and any other TCP/IP devices with built-in SNMP support. You should also know that SNMP can be used to configure devices remotely, monitor network performance, and detect network faults.

Though you don't necessarily need to know how to set it up, SNMP works through management systems and SNMP agents. Management systems monitor agents and determine if they have any system events that need to be reported to the administrative platform that is controlling them.

Network Monitor

Network Monitor is a downloadable tool available through Microsoft. The Monitor is a port and protocol analyzer designed to view and analyze network performance data. Sometimes it's helpful for administrators to know what type of traffic is being passed on the network. For example, you may not realize that you have a large amount of FTP traffic going to and from a particular server. It could mean that the server is being exploited for personal use by an employee.

Launching the Network Monitor will create a series of parsers (programs that search text via tokens) that will allow you to capture data. To capture data, you will need to select your network card, as shown in the figure below, and then click **New Capture**.

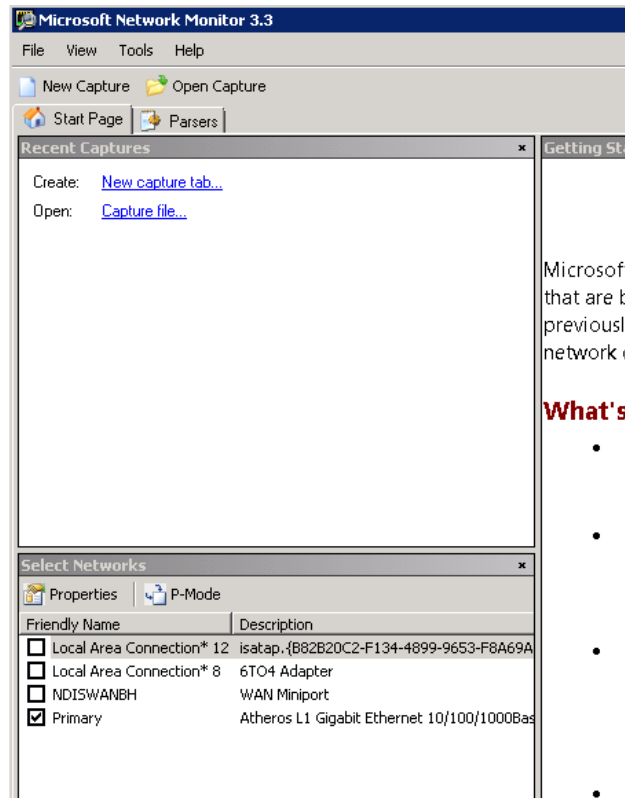


Figure 87: Starting a New Network Monitor Capture

Once you click **Capture**, this will launch the Network Capture Wizard. You will then need to click **Start**. At this point, network traffic will start to flow through the "All Traffic" monitor. You will then be able to see which programs are creating TCP/IP traffic. You'll be able to see which programs are creating the most traffic, and you can also look for signs of a network intrusion. Once you have captured the network data you are after, you can export these data by choosing the **Save As** button. This will export the data as a CAP file that you can open with the Network Monitor at a later time.

The screenshot displays the Microsoft Network Monitor 3.4 application window. The interface is divided into several panes:

- Network Conversations:** A tree view on the left showing traffic categories like 'All Traffic', 'My Traffic', and 'Other Traffic'.
- Display Filter:** A pane for applying filters to the captured data.
- Frame Summary:** A table listing captured frames with columns for Frame Number, Time Date Local Adjusted, Time Offset, Process Name, Source, Destination, Protocol Name, and Description. Frame 42 is highlighted.
- Frame Details:** A pane showing the structure of the selected frame (Frame 42), including Ethernet, IPv4, UDP, and HTTP details.
- Hex Details:** A pane showing the raw hexadecimal data of the selected frame.

At the bottom of the window, a status bar indicates: Version 3.4.2350.0, Displayed: 58, Dropped: 0, Captured: 58, Pending: 0, Focused: 42, Selected: 1.

Figure 88: Network Monitor

Another component of the Network Monitor package is a command-line program called "nmcap.exe". Among the several advantages of this program is that it's lightweight and won't slow your server down while it's capturing data. It's highly configurable, and can be used in scripts or batch files. It can also be set to run automatically. When the data capture is done, you'll be able to view the capture data in the graphical Network Monitor utility.

Practice Questions

Chapter 1

1. Your organization consists of a single Windows Server 2008 Active Directory domain that is spread across two IP subnets. One DHCP server exists in the domain. You receive a complaint from a user that she is unable to access domain resources. You discover that her computer's IP address is 169.254.11.213. What action should you perform?

 - A. Install a DHCP Relay Agent on the user's local subnet.
 - B. Install a WDS server on the user's local subnet.
 - C. Enable traffic on TCP 1542 as a Windows Firewall exception on the user's computer.
 - D. Disable PXE on the user's computer.
2. Your organization consists of a multi-domain Active Directory forest in which all servers run Windows Server 2008 and all client computers run Windows Vista. Hosts in all domains are configured to use both IPv4 and IPv6. You need to ensure that hosts in different domains in the forest can communicate through networks that are separated by NAT firewalls. Your solution must involve the least amount of administrative effort. What action should you perform?

 - A. Enable the Peer Name Resolution Protocol (PNRP) on all organizational firewalls.
 - B. Enable Teredo on all organizational firewalls.
 - C. Create static NAT translation entries on all organizational firewalls.
 - D. Create at least one IPv6 scope on all DHCP servers.
3. Your organization includes a newly deployed Windows Server 2008 e-mail server that is configured with the following IP configuration information:

IP Address: 172.16.16.3
Subnet mask: 255.255.224.0

You receive complaints from users who state that they are unable to connect to the e-mail server. What action should you perform?

 - A. Change the server's IP address to 172.16.0.3.
 - B. Change the server's IP address to 172.16.32.3.
 - C. Configure the server to use a /20 subnet mask.
 - D. Configure the server to use the subnet mask 255.255.128.0.

4. Your organization consists of an IP internetwork that is routed by a multihomed Windows Server 2008 member server that is configured with the RRAS server role. You need to configure a persistent default route on the server from the command prompt that sends all default traffic out of the interface with IP address 192.168.1.1. What action should you perform?
- A. Issue the command `route print 192.168.1.0` on the server.
 - B. Issue the command `route -persistent 192.168.1.0` on the server.
 - C. Issue the command `route -p add 255.255.255.255 mask 255.255.255.255 192.168.1.1` on the server.
 - D. Issue the command `route -p add 0.0.0.0 mask 0.0.0.0 192.168.1.1` on the server.
5. Your organization consists of a single Active Directory domain in which all servers run Windows Server 2008 and all client computers run Windows Vista. You are planning to migrate from IPv4 to IPv6 and need to educate yourself as to how IPv6 addressing works. Which of the following represents an invalid IPv6 address?
- A. 2001:cdba:0000:0000:0000:0000:3257:9656
 - B. 2001:cdba:0:0:0:0:3257:9655
 - C. 2001:cdba::3257:9651
 - D. 2001:cgba::3257:9654
6. Your organization is planning to migrate from an IPv4 infrastructure to an IPv6 infrastructure. Your manager is concerned about how IPv6 packets can be routed over the public Internet, especially to destinations that still use IPv4. What actions should you perform? (Select two choices. Each correct answer represents an independent solution.)
- A. Deploy the Teredo transition technology in your network.
 - B. Deploy NAT in your network.
 - C. Deploy NPS in your network.
 - D. Deploy 6to4 technology in your network.

Chapter 2

1. Your organization has a single Windows Server 2008 Active Directory domain named Birdco.com. Your DNS infrastructure includes both Windows Server 2008 DNS servers and Linux servers running the latest version of BIND. You need to insure that DNS zone data is not copied by unauthorized individuals. What action should you perform?
- A. Configure the zones on the BIND DNS servers to integrate with Active Directory.
 - B. Specify the Allow zone transfers only to the following servers option on your Windows Server 2008 DNS servers.
 - C. Disable zone transfers on your Windows Server 2008 DNS servers.
 - D. Configure secondary notification on your Windows Server 2008 DNS servers.

2. Your organization has a single Active Directory forest with an Active Directory domain in New York named Birco.com, a domain in Canada named Canada.Birdco.com, and a domain in London named London.Birdco.com. Users in the London domain complain of poor DNS query performance when they attempt to resolve host names from the Canada and root domains. You need to maximize name resolution performance for this group of users. Your solution must involve the least amount of administrative effort. What action should you perform?
- A. Deploy a HOSTS file to all computers in the London domain.
 - B. Configure connection-specific DNS suffixes for all hosts in the London domain.
 - C. Deploy a GlobalNames zone in all Active Directory domains.
 - D. Deploy a DNS Suffix Search List via a GPO for the London users.
3. Your organization has a single Windows Server 2008 Active Directory domain. Two servers are configured with the DNS Server role. These DNS servers host a single Active Directory-integrated zone. You change the IP address of a network printer on the network, and you now need to remove the inverse lookup record for the printer immediately. What action should you perform?
- A. Run the command `dnscmd /zoneddelete` on either DNS server.
 - B. Run the command `dnscmd /zoneddelete` on both DNS servers.
 - C. Run the command `dnscmd /recorddelete` on either DNS server.
 - D. Run the command `dnscmd /recorddelete` on both DNS servers.
4. Your organization consists of a single Active Directory forest organized as a single Windows Server 2008 domain. The domain includes four servers that are configured with the DNS Server role. The DNS servers host a single Active Directory-integrated zone. You plan to decommission WINS in the very near future, yet you need to maintain single label name resolution for key servers. You have run the command `dnscmd /enableglobalnamesupport` on all DNS servers and have created a zone named GlobalNames. What action should you perform next?
- A. Add NS records to the GlobalNames zone for the key servers.
 - B. Add CNAME records to the GlobalNames zone for the key servers.
 - C. Run the command `dnscmd /zoneadd` on a DNS server.
 - D. Create at least one IPv6 scope in the forest.

5. Your organization has a main office and a branch office. The main office contains a Windows Server 2008 domain controller named MAIN01 that hosts a standard primary DNS zone. The branch office contains a domain controller named BRANCH01 that hosts a standard secondary zone. Because you have moved some infrastructure servers in the main office to a different IP subnet, you need to ensure that these resource record changes are propagated immediately to the branch office. You need to be mindful of traffic on the WAN connection that links the two offices. What action should you perform?
- A. Issue the command `dnscmd /zonereoad` on MAIN01.
 - B. Issue the command `dnscmd /zonerefresh` on MAIN01.
 - C. Issue the command `dnscmd /zonereoad` on BRANCH01.
 - D. Issue the command `dnscmd /zonerefresh` on BRANCH01.
6. You have upgraded all servers in your company's single Active Directory domain to Windows Server 2008. The organization's SharePoint Server intranet Web site is hosted on four different Web servers. The web servers are configured with identical hardware and each one is assigned a unique IP address. You need to configure the network such that incoming SharePoint connection requests are distributed evenly across all four Web servers. What action should you perform?
- A. Enable netmask ordering at the DNS zone level.
 - B. Enable round-robin at the DNS zone level.
 - C. Enable netmask ordering at the DNS server level.
 - D. Enable round-robin at the DNS server level.
7. Your network consists of a multi-domain Active Directory forest in which all servers run Windows Server 2008 and all client computers run Windows Vista. Your company recently entered into a strategic partnership with another company who operates a Windows Server 2008 forest. You have configured a forest trust with the partner, and now want to maximize the efficiency of name resolution for your users, many of whom will access resources that are located in the remote forest. Your solution must involve the least amount of administrative effort. What action should you perform?
- A. Deploy HOSTS files to all client computers.
 - B. Instruct users to add the appropriate DNS suffixes from the remote forest to their network adapter TCP/IP properties.
 - C. Deploy a custom DNS suffix search list to users via a Group Policy Object (GPO).
 - D. Disable recursion on your local DNS servers.

Chapter 3

1. Your organization consists of a single Active Directory domain in which all servers run Windows Server 2008 and all client computers run Windows Vista with Service Pack 1. You plan to install a Windows Server 2008 VPN server named SECURERA01. Corporate security policy states that users' logon credentials not be passed between their remote device and the VPN server. Mutual authentication must also occur between the VPN client and the VPN server. What action should you perform?
 - A. Configure PPTP tunneling on SECURERA01.
 - B. Configure L2TP tunneling on SECURERA01.
 - C. Configure SSTP tunneling on SECURERA01.
 - D. Configure SSL tunneling on SECUREA01.

2. Your network consists of a single Active Directory domain in which all servers run Windows Server 2008. You are planning a secure remote access infrastructure that includes three servers: WINNPS: Network Policy Server, HEALTH01: System Health Validation Server, Remediation Server, VPN01: VPN Server. You need to ensure that VPN client computers are screened by network health policies. What action should you perform to complete the configuration?
 - A. Configure VPN01 as a System Health Validator.
 - B. Configure VPN01 as a RADIUS server.
 - C. Configure VPN01 as a RADIUS client of HEALTH01.
 - D. Configure VPN01 as a RADIUS client of WINNPS.

3. You manage an Active Directory domain named Birdco.com in which all servers run Windows Server 2008 and all client computers run Windows Vista. Half of the client computers are laptop computers in which users associate with local wireless access points for their connectivity. You are in the process of deploying 802.1X NAP enforcement on the network. You need to configure all wireless clients for health policy monitoring and remediation. What actions should you perform? (Select two answers. Each correct choice represents a part of a single solution.)
 - A. Deploy an IEEE 802.3 Group Policy.
 - B. Deploy an 802.11 Group Policy.
 - C. Deploy a Group Policy that configures wireless client computers as RADIUS clients.
 - D. Deploy a Group Policy that enables the Windows Security Center.

4. You need to secure Remote Desktop sessions between your company's Windows Vista-based administrative workstations and your company's Windows Server 2008-based domain controller. What actions should you perform? (Select two choices. Each correct choice represents a part of a single solution.)
- A. Enable a security layer for RDP connections in Group Policy.
 - B. Disable Single Sign-On for Terminal Services.
 - C. Ensure that terminal servers are placed in the same Active Directory OU as other servers on the network.
 - D. Change the default RDP port on the server and the client computers.
5. You need to secure Remote Desktop connections between 20 Windows Vista-based administrative workstations and 5 Windows Server 2008-based domain controllers. You plan to deploy a Group Policy Object (GPO) that defines the use of an RDP security layer. Your solution must require that the Terminal Server authenticate itself to all RDP clients. What action should you perform?
- A. Deploy the Negotiate security layer for RDP connections.
 - B. Deploy the RDP security layer for RDP connections.
 - C. Deploy the SSL (TLS 1.0) security layer for RDP connections.
 - D. Require secure RPC communication between the Terminal Server and the client computers.
6. Your network consists of Windows Server 2008 servers and Windows Vista client computers. You have configured a domain member server named VPN01 as a VPN server. The VPN server is associated with a NPS server named POLICY01 that is set up for NAP health policy enforcement. Users complain that they are unable to access resources that lie beyond the VPN server when their connection is authenticated. You have verified that the security health validator is functional and that the remote computers are in compliance with health policy. What action should you perform?
- A. Change the network access permission for all user accounts to Allow access.
 - B. Open the appropriate ports on the corporate firewall.
 - C. Install Service Pack 1 on all Windows Vista computers.
 - D. Verify the IP address pools on the VPN server.

7. You are planning to upgrade a Windows Server 2003 Routing and Remote Access (RRAS) server to Windows Server 2008. Your manager asks if any RRAS protocols have been removed in Windows Server 2008. You need to know this information to prevent a loss of functionality post-upgrade. Which of the following remote access protocols have been removed in Windows Server 2008 RRAS? (Select two answers.)
- A. OSPF
 - B. BAP
 - C. RIPv2
 - D. CHAP

Chapter 4

1. Your organization consists of a single Active Directory domain named Birdco.com in which all servers run Windows Server 2008. The printing infrastructure consists of 20 print queues that are located in the main office, and 10 print queues that are located in the branch office. You open Print Management console on a member server in the main office and load up all print servers across the entire domain. In order to create a toner resupply schedule, you need to gauge how busy each printer is over the course of each workday. Your solution must involve the least amount of administrative effort. What action should you perform?
- A. Configure e-mail notification on the Printers Not Ready filter in the Print Management console.
 - B. Configure e-mail notification on the Printers With Jobs filter in the Print Management console.
 - C. Ensure all printers are deployed to Active Directory.
 - D. Migrate all print queues to the server on which you have started the Print Management console.
2. You manage a Windows Server 2008 file server named FILESERV01 that hosts a domain DFS namespace named \birdco.com\docs. You need to ensure that domain users can add content to a DFS folder named REPORTS. However, your solution must adhere to the IT security principle of least privilege and must involve the least amount of administrative effort. What action should you perform? (Select two choices. Each correct choice represents a part of a single solution).
- A. Set the share permissions for the REPORTS folder targets such that the Administrators group is granted the Owner permission level.
 - B. Set the NTFS permissions for the REPORTS folder targets such that the Administrators group is granted the Traverse Folder permission.
 - C. Set the share permissions for the REPORTS folder targets such that the Domain Users group is granted the Contributor permission level.
 - D. Set the share permissions for the REPORTS folder targets such that the Domain Users group is granted the Co-Owner permission level.

3. You manage a Windows Server 2008 member server named FS01 on which the File Services role is installed. The server contains a shared folder named MEDIA and users can store project files inside the MEDIA folder. You need to ensure that users are unable to save peer-to-peer file sharing connection files on the MEDIA share. What action should you perform? (Select two choices. Each correct choice represents a part of a single solution).
- A. Apply the quota template to the MEDIA folder.
 - B. Apply the file screen template to the MEDIA folder.
 - C. Create a new file group and file screen template.
 - D. Create a new quota template.

Chapter 5

1. Your company consists of a single Active Directory domain in which all servers run Windows Server 2008. You manage a Windows Server 2008 application servers named MOSS01. Users complain of receiving extremely slow connection times to MOSS01. You need to view real-time network performance data to isolate the cause of the problem. What action should you perform?
- A. View networking counters in System Monitor.
 - B. Examine the Network Resource Overview in Reliability and Performance Monitor.
 - C. Analyze the Performance tab of Task Manager.
 - D. Create an event subscription for network-related events on MOSS01.
2. You need to capture network traffic data on a Windows Server 2008 member server named DIOGENES. However, you are connected to DIOGENES via a Terminal Services session and do not want to capture this traffic. What action should you perform?
- A. Run the command `nmcap /network * /capture /file test.cap`.
 - B. Run the command `nmcap /network * /capture "TS" /StopWhen /TimeAfter 2 min /file TS.cap`.
 - C. Run the command `NMcap /network * /capture "!(tcp.port == 3389)" /file test.cap`.
 - D. Run the command `NMcap /network * /capture "!(tcp.port == 389)" /file test.cap`.

Answers & Explanations

Chapter 1

1. Answer: A

Explanation A. Correct. In Windows Server 2008, the DHCP Relay Agent service is installed as a service of the RRAS server role. The DHCP Relay Agent allows DHCP broadcast messages to cross router interfaces.

Explanation B. Incorrect. Windows Deployment Service (WDS) is an operating system deployment technology that relies upon DHCP. However, the WDS service itself will not allow DHCP Discover packets to cross router interfaces.

Explanation C. Incorrect. DHCP Relay is defined in Request for Comments (RFC) 1542. Traditional DHCP traffic is not blocked by Windows Firewall due to its mission-critical nature.

Explanation D. Incorrect. The Preboot Execution Environment (PXE) represents an alternate way for a device to obtain an IP address lease from DHCP. However, PXE does not factor into this scenario.

2. Answer: B

Explanation A. Incorrect. PNRP is an IPv6-based discover technology that underpins Microsoft applications such as Meeting Space. However, without Teredo enabled on all firewalls, the firewalls will not let IPv6 packets traverse NAT.

Explanation B. Correct. Teredo technology allows IPv6 packets to traverse IPv4 Network Address Translation (NAT) firewalls. The Teredo client is enabled by default in Windows Server 2008 and Windows Vista.

Explanation C. Incorrect. Although this would solve the problem, it is almost mind-numbingly complex compared to enabling Teredo on all firewalls.

Explanation D. Incorrect. This answer choice represents a classic “red herring.” That is, we don’t need to worry about IP addressing. There is nothing in the scenario that would lead us to believe there is a problem with addressing.

3. Answer: C

Explanation A. Incorrect. The problem here is that the given IP address and the subnet mask do not match.

Explanation B. Incorrect. The problem here is in the subnet mask.

Explanation C. Correct. The given IP address is invalid with a 19-bit subnet mask. Here we need to use a /20 (255.255.240.0 in decimal) subnet mask.

Explanation D. Incorrect. This subnet mask masks only 17 bits, which continues to render the server’s IP address as invalid.

4. Answer: D

Explanation A. Incorrect. The Route Print command simply displays all or part of the server's routing table.

Explanation B. Incorrect. The -p flag of the route command is used, not the word "persistent." Besides, the syntax of this statement is incorrect.

Explanation C. Incorrect. We need to use decimal 0, not decimal 255, in our route statement.

Explanation D. Correct. The -p flag of the route command makes the static route persistent across server reboots.

5. Answer: D

Explanation A. Incorrect. This is a valid IPv6 address.

Explanation B. Incorrect. This is a valid IPv6 address.

Explanation C. Incorrect. This is a valid IPv6 address.

Explanation D. Correct. This IPv6 address is invalid because g is an illegal character in hexadecimal notation.

6. Answers: A, D

Explanation A. Correct. Teredo is a vendor-neutral technology that was developed to allow IPv6 and IPv4 to interoperate.

Explanation B. Incorrect. Actually, NAT will become a deprecated technology once the world has fully migrated to IPv6. NAT is relevant only with IPv4 addresses and was developed to prevent IPv4 address depletion.

Explanation C. Incorrect. The Windows Server 2008 Network Policy Server (NPS) technology is not relevant in this scenario.

Explanation D. Correct. The 6to4 technology was developed to facilitate communications between IPv6 hosts and IPv4 hosts.

Chapter 2

1. Answer: B

Explanation A. Incorrect. You cannot store non-Windows DNS zones in Active Directory. Only Windows DNS servers can store DNS zone data in an Active Directory partition.

Explanation B. Correct. You have great control over how Windows Server 2008 DNS servers that host primary zones participate in zone transfers. This is good news from an IT security perspective.

Explanation C. Incorrect. Because the scenario states that we use a mixed Windows/Linux DNS infrastructure, reason tells us that we need to perform zone transfers between these DNS servers.

Explanation D. Incorrect. The Notify option in Windows Server 2008 DNS configures the DNS server to notify any connected secondary DNS servers of zone database changes.

2. Answer: D

Explanation A. Incorrect. This technique is not a best practice and requires far too much administrative effort to deploy and maintain over time.

Explanation B. Incorrect. While this solution would work, it requires far more administrative effort than simply deploying a GPO with a DNS Suffix Search List.

Explanation C. Incorrect. The new GlobalNames zone simply provides single label name resolution in lieu of WINS in an Active Directory infrastructure. This action won't necessarily improve name resolution times across domains.

Explanation D. Correct. You can provide TCP/IP clients in a domain with a list of DNS suffixes to ease name resolution by deploying a Group Policy Object (GPO) with the DNS Suffix Search List policy enabled.

3. Answer: C

Explanation A. Incorrect. The zonedefile parameter is used to "nuke" an entire DNS zone file, which is not what we want to do here.

Explanation B. Incorrect. Not only do we not want to delete an entire zone in this case, but we certainly don't want to remove the zone on both DNS servers.

Explanation C. Correct. The recorddelete command can be used to delete any kind of resource record from a DNS zone. For instance, you would include the parameter PTR to delete an inverse lookup, or pointer, resource record.

Explanation D. Incorrect. Because the two servers host a single Active Directory-integrated zone, there is no need whatsoever to perform this action on both boxes. AD replication will propagate the change automatically.

4. Answer: B

Explanation A. Incorrect. Remember that NS records are used only to identify authoritative DNS servers.

Explanation B. Correct. The GlobalNames zone does not support dynamic client update; in other words, you must add the records yourself. These records can be host (A) or alias (CNAME) record types.

Explanation C. Incorrect. The scenario states that we have already created the standard primary zone named GlobalNames (globalnames is also supported; the name is not case-sensitive). AD replication will take care of propagating the zone to all DNS servers in the domain.

Explanation D. Incorrect. The GlobalNames DNS zone works with both IPv4 and IPv6; there is not a requirement for IPv6.

5. Answer: D

Explanation A. Incorrect. Not only do we need to initiate the zone transfer from the branch office and not the main office, we need to conserve WAN bandwidth, which performing a full zone reload will not do.

Explanation B. Incorrect. We do want to perform a zone refresh (incremental zone transfer); however, we need to initiate this action from BRANCH01, not MAIN01.

Explanation C. Incorrect. We need to conserve WAN bandwidth in this case, so pulling the entire zone file from MAIN01 to BRANCH01 is inefficient.

Explanation D. Correct. We want to force an incremental zone transfer from MAIN01 to BRANCH01.

6. Answers: B, D

Explanation A. Incorrect. Not only do we not need to enable netmask ordering here, but the option is available only at the DNS server level, not the zone level.

Explanation B. Correct. While we do need to enable round-robin in this case, we do so at the DNS server level and not at the zone level.

Explanation C. Incorrect. Netmask ordering makes name resolution more efficient; the resource record that corresponds to the nearest host is returned to the client when more than one record exists for the same service. However, in this case we are concerned with round-robin, not netmask ordering.

Explanation D. Correct. Round-robin is a feature of Windows Server DNS that is enabled at the server (service) level. It distributes incoming requests for service equitably among multiple authoritative DNS servers.

7. Answer: C

Explanation A. Incorrect. This solution is inflexible and requires a staggering amount of administrative effort to implement and maintain.

Explanation B. Incorrect. Firstly, it is hoped that ordinary users don't have administrative privileges on their computers. Secondly, this "solution" requires far too much administrative effort.

Explanation C. Correct. Windows Server 2008 contains a Group Policy setting that enables users to receive a list of DNS suffixes that will be appended automatically to unqualified DNS names in their queries.

Explanation D. Incorrect. If we disable recursion on our DNS servers, the servers will never resolve names for which the servers are not personally authoritative.

Chapter 3

1. Answer: B

Explanation A. Incorrect. Point to Point Tunneling Protocol (PPTP) does not support mutual authentication. What is worse is that PPTP sends users' credentials over the "pipe" before the secure channel is established between the VPN client and the VPN server.

Explanation B. Correct. L2TP tunneling supports integrity, confidentiality, and mutual authentication when it is deployed with the Internet Protocol Security (IPSec) security protocol.

Explanation C. Incorrect. Secure Socket Tunneling Protocol enables secure VPN connections across firewalls that block L2TP/IPSec ports. SSTP prevents user credentials from crossing the connection until after a socket has occurred. However, SSTP requires other protocols to support mutual authentication.

Explanation D. Incorrect. We cannot create a VPN tunnel simply with Secure Sockets Layer (SSL). However, SSTP uses SSL in creating secure remote access connections.

2. Answer: D

Explanation A. Incorrect. We already have a system health validator in the mix. In this scenario we simply need to point the VPN server to it.

Explanation B. Incorrect. In order to link the VPN server to our Health Validation and NPS system, we need to configure VPN01 as a RADIUS client to the NPS RADIUS server.

Explanation C. Incorrect. In this case the RADIUS server is our NPS device. Remember that Network Policy Services is the Windows Server 2008 replacement for the Internet Authentication Service (IAS) that we knew and loved in Windows Server 2003.

Explanation D. Correct. In order to link our VPN server to our Health Validation server, we need to point the VPN server (as a RADIUS client) to our RADIUS server which, in this case, is WINNPS.

3. Answers: B, D

Explanation A. Incorrect. We actually need to deploy wireless connection settings by using the Wireless Network (IEEE 802.11) Policies in Windows Server 2008 Group Policy.

Explanation B. Correct. The 802.11 Group Policy settings allow you to specify which WLANs a user is or is not allowed to associate with and the security settings that govern each allowed connection.

Explanation C. Incorrect. In this scenario we need to configure our wireless access points (WAPs) as RADIUS clients of our Network Policy Server (NPS) server.

Explanation D. Correct. The Windows Security Health Validator service requires that the Security Center be running on all Windows Vista clients who need to be part of a network health policy.

4. Answers: A, D

Explanation A. Correct. In Windows Server 2008, you can enable one of three security layers for RDP connections with Windows Vista or other Windows Server 2008 computers: Negotiate, RDP, and SSL (TLS 1.0).

Explanation B. Incorrect. Actually, Microsoft recommends enabling SSO for Terminal Services to reduce the exposure of user credentials over the network.

Explanation C. Incorrect. Microsoft recommends that you place Terminal Server computers in their own OU to better scope policy.

Explanation D. Correct. Although this is a drastic step, it is a legitimate way to harden Terminal Services in certain high-risk situations.

5. Answer: C

Explanation A. Incorrect. Negotiate policies can use either RDP encryption or SSL (TLS 1.0) encryption, but it is possible that the server is not required to authenticate if RDP encryption is negotiated.

Explanation B. Incorrect. The RDP security level does encrypt the Remote Desktop Protocol (RDP) traffic. However, the server is not authenticated in this situation.

Explanation C. Correct. When you use Secure Sockets Layer (SSL) encryption for the RDP connection, the server authenticates itself to the client computer. This provides higher security (although with more administrative overhead) than RDP encryption.

Explanation D. Incorrect. We are securing the transport protocol (RDP) in this scenario, not RPCs.

6. Answer: D

Explanation A. Incorrect. Because we are using NAP/NPS, we should ensure that all domain user accounts with remote access aspirations are configured with the option Control access through NPS.

Explanation B. Incorrect. The users in this scenario are being authenticated which tells us that the problem has nothing to do with firewall configuration.

Explanation C. Incorrect. Windows Vista RTM ships with the NAP client. You might recall that only Service Pack 3 provides the NAP client to Windows XP Professional computers.

Explanation D. Correct. In all likelihood, there is a borked connection between the VPN server and the domain's DHCP server. Alternatively, the VPN server could be configured to dole out invalid IP addresses from a static pool.

7. Answers: A, B

Explanation A. Correct. The Open Shortest Path First (OSPF) routing protocol has been removed in Windows Server 2008 RRAS.

Explanation B. Correct. Bandwidth Allocation Protocol (BAP) is not available in Windows Server 2008 RRAS.

Explanation C. Incorrect. Routing Information Protocol version 2 (RIPv2) is alive and well in Windows Server 2008 RRAS.

Explanation D. Incorrect. The industry standard Challenge Handshake Authentication Protocol (CHAP) is very much present in Windows Server 2008 RRAS.

Chapter 4

1. Answer: B

Explanation A. Incorrect. You can indeed create display filters in the Print Management console, and these filters can send e-mail messages when new items are added. However, we want to activate the built-in Printers With Jobs filter to gauge how busy each printer is over the course of each workday.

Explanation B. Correct. By enabling e-mail notification on this built-in printer filter, we can get a handle on which print queues are most active and plan our toner resupply schedule accordingly.

Explanation C. Incorrect. Whereas it is true that the Print Management console allows you to easily deploy print queues to Active Directory via Group Policy, this does not help us meet our goal.

Explanation D. Incorrect. The Printer Migration feature in Windows Server 2008 involves exporting and importing print queues across servers. There is nothing in the scenario that leads us to conclude that we want to change the location of any of our print queues and/or print servers.

2. Answers: A, C

Explanation A. Correct. Good Windows administrative practice dictates that administrators have full control over all shared folders in a DFS folder hierarchy.

Explanation B. Incorrect. We don't need to worry about NTFS permissions here so much as we do the share permissions which make the DFS folder available on our network.

Explanation C. Correct. The Contributor permission level is "just right" in this case. We might want to tweak up the permissions a bit more by using Authenticated Users instead of Domain Users, but this is fine for our purposes here.

Explanation D. Incorrect. We need to take care not to over-privilege our users. By granting the Domain Users (or, better yet, the Authenticated Users) group the Contributor permission level, these users can add or remove content to the folder targets without having too much permission.

3. Answers: B, C

Explanation A. Incorrect. We are concerned with file screens in this scenario, not quota templates.

Explanation B. Correct. After we define the file screen template, the final step in this configuration is to define a new file screen that attaches our custom file screen template.

Explanation C. Correct. The first thing we would do is to create a new file group that registers the appropriate file extensions. Next, we would create a reusable file screen template that attaches the new file group.

Explanation D. Incorrect. We need to perform file screening here, not simply cap the maximum amount of server disk space users can access.

Chapter 5

1. Answer: B

Explanation A. Incorrect. The new Reliability and Performance Monitor tool in Windows Vista and Windows Server 2008 fully takes the place of the System Monitor tool that was present in earlier versions of Windows.

Explanation B. Correct. The Resource Overview section of the Reliability and Performance Monitor tool gives real-time data on the CPU, Disk, Memory, and Network subsystems of a local or remote computer.

Explanation C. Incorrect. You could use Task Manager to see real-time networking statistics; however, to do this you need the Networking tab, not the Performance tab.

Explanation D. Incorrect. In this case we need real-time networking statistics, not logged data. Besides, Event Viewer and event log subscriptions give just information, warning, and error state data.

2. Answer: C

Explanation A. Incorrect. This command starts an unfiltered network capture from the command line. Therefore, the captured data would include Terminal Services streams as well.

Explanation B. Incorrect. This capture syntax is more explicit than the one given in answer choice 1. However, we still are not filtering out Terminal Services (TS) traffic.

Explanation C. Correct. This syntax will capture traffic on DIOGENES and exclude Terminal Services traffic which operates on TCP port 3389 by default.

Explanation D. Incorrect. TCP port 389 maps to the well-known port number for the Lightweight Directory Access Protocol (LDAP), not Terminal Services Remote Desktop Protocol (RDP) traffic.