

Microsoft
Server 2003
Planning Network Infrastructure
(70-293)

Microsoft Certified
Systems Engineer (MCSE)

**Smarter
Training**

This LearnSmart exam manual is designed to take you to the next level of success by helping you prepare for the Server 2003 Planning Network Infrastructure exam (70-293). By studying this exam, you will become familiar with an array of exam-related topics, including:

- Server Roles and Server Security
- Network Infrastructure
- Routing and Remote Access
- Server Availability
- Network Security
- And more!

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

Server 2003 Planning Network Infrastructure (70-293) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC
Product ID: 10013
Production Date: July 12, 2011
Total Questions: 25

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

Volume, Corporate, and Educational Sales	2
Introduction	7
What to Know	7
Tips	7
Resources	8
Planning and Implementing Server Roles and Server Security	9
Configure security for servers that are assigned specific roles	9
Plan a secure baseline installation	10
<i>Plan a strategy to enforce system default security settings on new systems</i>	<i>11</i>
<i>Identify client operating system default security settings</i>	<i>11</i>
<i>Identify all server operating system default security settings.....</i>	<i>11</i>
Plan security for servers that are assigned specific roles. Roles might include domain controllers, Web servers, database servers, and mail servers	12
<i>Deploy the security configuration for servers that are assigned specific roles.....</i>	<i>13</i>
<i>Create custom security templates based on server roles</i>	<i>13</i>
Evaluate and select the operating system to install on computers in an enterprise	15
<i>Identify the minimum configuration to satisfy security requirements</i>	<i>16</i>
Planning, Implementing, and Maintaining a Network Infrastructure.....	16
Plan a TCP/IP network infrastructure strategy	16
<i>Analyze IP addressing requirements</i>	<i>18</i>
<i>Plan an IP routing solution</i>	<i>19</i>
<i>Create an IP subnet scheme</i>	<i>20</i>
<i>CIDR and VLSM.....</i>	<i>20</i>
Plan and modify a network topology	21
<i>Plan the physical placement of network resources</i>	<i>22</i>
<i>Identify network protocols to be used</i>	<i>23</i>
Plan an Internet connectivity strategy	24
Plan network traffic monitoring. Tools might include Network Monitor and System Monitor.....	25
<i>Capturing Data</i>	<i>27</i>
Troubleshoot connectivity to the Internet	27
<i>Diagnose and resolve issues related to Network Address Translation (NAT)</i>	<i>27</i>
<i>Diagnose and resolve issues related to name resolution cache information</i>	<i>27</i>

<i>Diagnose and resolve issues related to client configuration</i>	28
Troubleshoot TCP/IP addressing	28
<i>Diagnose and resolve issues related to DHCP server address assignment</i>	28
Plan a host name resolution strategy	29
Plan a DNS namespace design	30
<i>Plan zone replication requirements</i>	30
<i>Plan a forwarding configuration</i>	32
<i>Plan for DNS security</i>	33
<i>Examine the interoperability of DNS with third-party DNS solutions</i>	34
Plan a NetBIOS name resolution strategy	35
<i>Plan a WINS replication strategy</i>	37
<i>Plan NetBIOS name resolution by using the Lmhosts file</i>	38
Troubleshoot host name resolution	39
<i>Diagnose and resolve issues related to DNS services</i>	39
<i>Diagnose and resolve issues related to client computer configuration</i>	39
Planning, Implementing, and Maintaining Routing and Remote Access	40
Plan a routing strategy	40
<i>Identify routing protocols to use in a specified environment</i>	40
RIP	42
RIP Version 1	42
RIP Properties	43
RIP Version 1 Shortcomings	43
RIP Version 2	44
OSPF	44
<i>Plan routing for IP multicast traffic</i>	45
Plan security for remote access users	46
<i>Plan remote access policies</i>	46
<i>Analyze protocol security requirements</i>	46
<i>Plan authentication methods for remote access clients</i>	47
Implement secure access between private networks	48
<i>Create and implement an IPSec policy</i>	48
Troubleshoot TCP/IP routing; Tools might include the route, tracert, ping, pathping, and netsh commands and Network Monitor	49

Planning, Implementing, and Maintaining Server Availability	50
Plan services for high availability	50
<i>Plan a high availability solution that uses clustering services</i>	50
<i>Plan a high availability solution that uses Network Load Balancing</i>	50
Identify system bottlenecks by using System Monitor	52
Implement a cluster server	52
<i>Recover from cluster node failure</i>	53
Manage Network Load Balancing. Tools might include the Network Load Balancing Monitor, Microsoft Management Console (MMC) snap-in, and the WLBS cluster control utility	54
Plan a backup and recovery strategy	54
<i>Identify appropriate backup types. Methods include full, incremental, and differential</i> ...	55
<i>Plan a backup strategy that uses volume shadow copy</i>	55
<i>Plan system recovery that uses Automated System Recovery (ASR)</i>	57
Planning and Maintaining Network Security	58
<i>Configure network protocol security</i>	58
<i>Configure protocol security in a heterogeneous client computer environment</i>	58
<i>Configure protocol security by using IPSec policies</i>	59
Configure security for data transmission	60
<i>Configure IPSec Policy Setting</i>	60
Plan for network protocol security	61
<i>Specify the required ports and protocols for specified services</i>	61
<i>Plan an IPSec policy for secure network communications</i>	62
Plan secure network administration methods	62
<i>Create a plan to offer Remote Assistance to client computers</i>	63
<i>Plan for remote administration by using Terminal Services</i>	63
Plan security for wireless networks	64
<i>IEEE 802.11 Authentication</i>	64
<i>WEP</i>	65
Plan security for data transmission	67
<i>Secure data transmission between client computers to meet security requirements</i>	68
<i>Secure data transmission by using IPSec</i>	69
<i>Troubleshoot security for data transmission. Tools might include the IP Security Monitor MMC snap-in and the Resultant Set of Policy (RSOP) MMC snap-in</i>	70

Planning, Implementing, and Maintaining Security Infrastructure	72
Configure Active Directory service for certificate publication	72
Plan a public key infrastructure (PKI) that uses Certificate Services	73
<i>Identify the appropriate type of certificate</i>	
<i>authority to support certificate issuance requirements</i>	74
<i>Plan the enrollment and distribution of certificates</i>	75
<i>Plan for the use of smart cards for authentication</i>	76
Plan a framework for planning and implementing security	77
<i>Plan for security monitoring</i>	78
<i>Plan a change and configuration management framework for security</i>	79
Plan a security update infrastructure. Tools might include Microsoft Baseline Security Analyzer and Microsoft Software Update Services	79
Practice Questions	81
Answers and Explanations	98

Introduction

This Exam Manual will help you prepare for the Microsoft Planning and Maintaining a Microsoft Windows Server 2003 Network Infrastructure (70-293) exam. This exam covers topics such as planning, implementing and maintaining DNS, WINS, DHCP, RRAS, and IPSec. These are just a few components of a Windows Server 2003 Network Infrastructure.

What to Know

Microsoft's exam 70-293, "Planning and Maintaining a Microsoft Windows Server 2003 Network Infrastructure" is a core networking system requirement for the MCSE (Microsoft Certified Systems Engineer) on Microsoft Windows Server 2003 certification. It is designed for IT professionals who "work in the typically complex computing environment of medium to large companies." There is no specific prerequisite for this exam but Microsoft recommends that candidates have "at least one year of experience implementing and administering a network operating system" in the type of environment [described here](#). As well, Microsoft says that a candidate "should have at least one year of experience in the following areas: Implementing and administering a desktop operating system and Designing a network infrastructure."

"Beginning with the release of the Microsoft Windows Server™ 2003–related exams, Microsoft will provide numerical scores on exams. By the end of September 2003, most exams will have the new score report format, which includes this numerical score... The new scale requires a minimum passing score of 700 on all exams. The maximum score on the exams will vary depending on the complexity of the skills being measured." - [Microsoft Exam and Testing Procedures FAQ](#)

Note that Microsoft does not document the format of a particular exam: "Microsoft exams might include adaptive testing technology and simulation items. Microsoft does not identify the format in which exams are presented."

The topics covered by this exam include:

- Planning and Implementing Server Roles and Server Security
- Planning, Implementing, and Maintaining a Network Infrastructure
- Planning, Implementing, and Maintaining Routing and Remote Access
- Planning, Implementing, and Maintaining Server Availability
- Planning and Maintaining Network Security
- Planning, Implementing, and Maintaining Security Infrastructure

[Visit here for more information.](#)

Tips

You want to start with this Exam Manual. Allocate about two weeks of study time and dedicate at least an hour a night, preferably several hours. If you can get a hold of related LearnSmart's practice exams, then do it. They will get you comfortable with the test format Microsoft will test you on.

Resources

[CramSession MCSE 2003 Certification Page](#)

[CramSession Certification & Training MCSE 2003 Articles](#)

[MCSE 2003 Feedback and Discussion Board](#)

[Free Question of the Day](#)

[LearnSmart's Practice Exams](#)

Planning and Implementing Server Roles and Server Security

Configure security for servers that are assigned specific roles

In February of 2002 Microsoft halted all development projects until every project could be reviewed for security, and to make sure that every developer knew how to write secure code. That month had a huge impact on the development of Windows Server 2003. If you have played around with this operating system (which you should have by now), you will see that a lot of the old security issues have been taken care of. Microsoft is proud of this. Microsoft wants to make sure that system engineers know exactly what new and improved security features are available to secure Windows Server 2003 and how to use these new utility tools. You can be sure that Microsoft is going to test your knowledge on security issues on this exam in several areas that pertain to Windows Server 2003 network infrastructure.

Windows Server 2003 comes with two tools called the **Configure Your Server Wizard** and the **Manage Your Server** tool. Both will help you configure all the appropriate settings for functionality and security for that particular server role. One of the first things you will see after you install and reboot Windows Server 2003 is the Configure Your Server Wizard. Server roles include file, print, application, e-mail, remote access, DNS, DHCP, WINS and Terminal server. You will have the option to configure certain tasks or server roles on your Windows Server 2003 as shown in *Figure 1*.

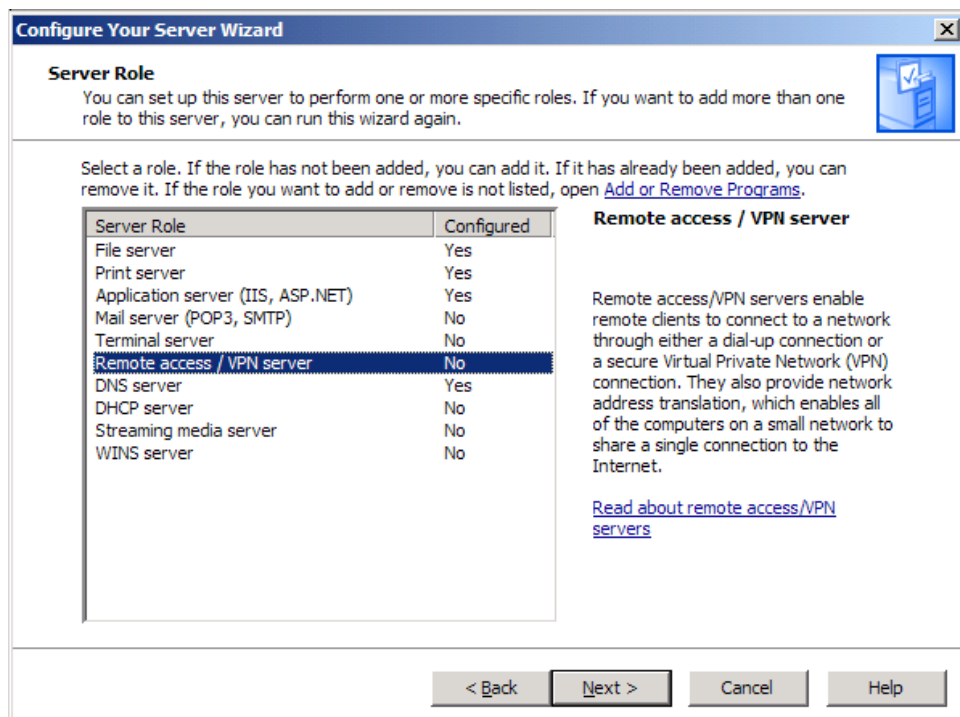


Figure 1 – Configure Your Server Wizard

Once you have installed and set up the server role you want, the Manage Your Server tool starts up as shown in *Figure 2*. The Manage Your Server tool can be used to add, remove, and manage server roles. It also lists the current roles that your server is currently performing along with providing tools to assist you in performing the role. For example, a file server role will have the option for setting up additional shares. The Manage Your Server utility also includes quick links to other options such as Administrative Tools, Windows Update page and Help and Support.

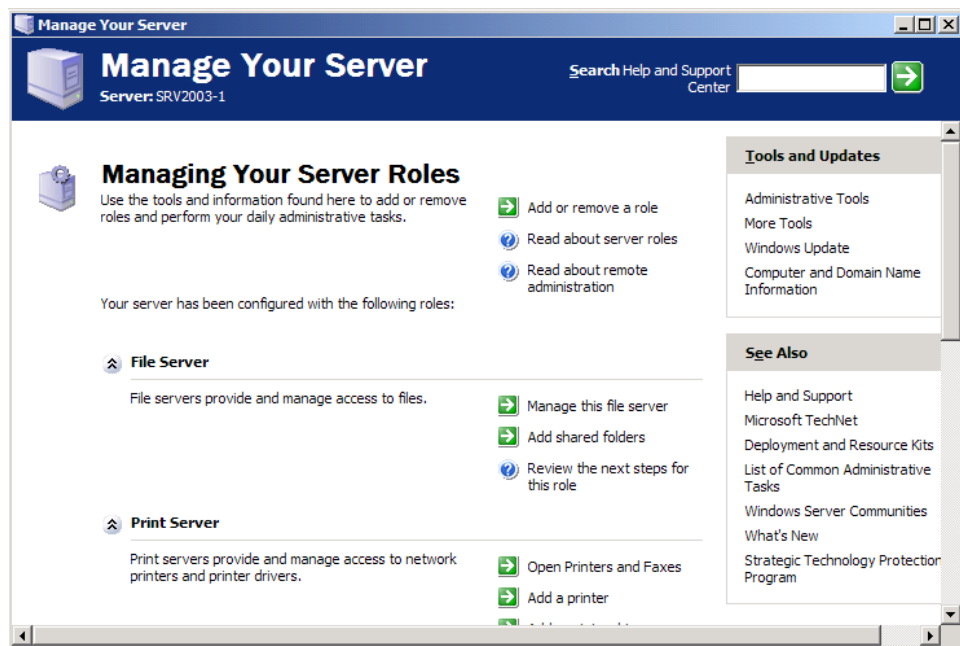


Figure 2 – Manage Your Server

Plan a secure baseline installation

Each new release of a Microsoft operating system has additional security enhancements. If your network happens to be one of a few that uses the latest server and client operating systems, then you already have a higher baseline of security than with networks that contain older operating systems. Advancements in technology have paved the way for increased security in regards to authentication, encryption, firewalls, etc. The catch is that these often work with their “partners,” meaning that the client has to have the latest technology in order to take advantage of the advanced security capabilities of the server. Having said that, there are also some practices that, regardless of your type of client operating systems, you can follow to enhance security.

Your strategy for hardening client operating systems should include the following:

- Where possible, upgrade to the latest operating system
- Install and continue to update the latest service packs, security patches and fixes on all clients and servers
- Use **Microsoft Baseline Security Analyzer (MBSA)** to scan clients and determine security weaknesses
- Uninstall or disable all unused services or protocols
- Use the host firewalls and filters built-in to the newer clients
- Use anti-virus software and update the anti-virus definition files regularly
- Use **Encrypted File System (EFS)** to secure sensitive information on company laptop computers

Plan a strategy to enforce system default security settings on new systems

The best way to ensure that your clients and servers are securely locked down would be to take advantage of **Active Directory's Group Policy** feature. This feature gives you the ability to apply settings and security lockdown features of the operating system. Not all operating systems support Group Policy: only Windows 2000, XP, and Server 2003 do. Each of those releases contains new features that can be applied through Group Policy. Make sure you are aware of what policies apply to what operating systems. You will want to initially use one of the built-in **Administrative Templates** when locking down default security settings on your computer systems.

Identify client operating system default security settings

Many networks have many different types of clients. You may have Windows 9X, NT, 2000, or even XP clients. Each one of those client operating systems has its own level of default security. With each new operating system (OS) release the default security improves. You need to review each type of operating system and identify the default security settings. For Windows 2000 and XP Professional clients, try to use group policies when configuring your client's default security settings. For clients that are older than Windows 2000 you can still use **System Policy** to secure portions of their operating systems. However, securing systems will be much easier when all of your clients are running at least Windows 2000 or above.

Identify all server operating system default security settings

Much like you client operating systems you may have several different server operating systems too. If you can, try to upgrade your servers up to at least Windows 2000, to take advantage of additional security features that are not offered in Windows NT. One feature that is of critical importance would be the ability to use and manage Remote Access Policies. Of course, that is only a concern for Domain Controllers. Those are not the only types of servers that you may have, but you will want to standardize some level of baseline security settings for all servers and then apply additional and more specific security settings for servers that have specific roles.

Plan security for servers that are assigned specific roles. Roles might include domain controllers, Web servers, database servers, and mail servers

Windows Server 2003 has the ability to be more than just a single purpose server. There are many different types of servers that you can configure with Windows Server 2003. Microsoft has made it very easy to not only configure and set up these server roles, but also to manage them.

Server Role	Description
Domain Controller	Stores directory data and is responsible for logon process, authentication, and directory services. By default all Domain Controllers exist in an OU named Domain Controllers and should have their security applied through Group Policy linked to that OU.
DNS Server	Used to resolve user-friendly domain names into IP addresses for communication on a network. Also has the ability to resolve services on the network to particular servers by using a resource record called SRV. For added security, you should use Active Directory Integrated DNS zones (secure updates) whenever possible.
DHCP Server	DHCP servers reduce administrative effort by automatically assigning from a defined scope an appropriate IP address and other communication IP addresses to clients. Hackers like to gain access to this server to obtain information so that they can spoof your network. You should keep these servers secure by using NTFS volumes and by limiting the number of administrators who have access to this server.
WINS Server	Used to resolve NetBIOS names into IP addresses for down-level clients, servers and applications. Down-level clients include Windows NT Workstation and Windows 9.x clients. As soon as you have no clients, servers or applications using NetBIOS you should uninstall this service.
Terminal Server	Terminal Servers (TS) allows you to remotely administer servers and run applications on the TS remotely. For maximum security when using TS, be sure to use the highest level of encryption.
Mail Server	Mail servers are used to implement POP3 (inbound mail) and SMTP (outbound mail) services. To secure mail servers use NTFS volumes and also use S/MIME or PGP to secure the communication to and from this server if needed.
File Server	File Servers store files, folders and applications accessed via the network through shared folders. NTFS volumes must be used to ensure that only the appropriate users or groups have access to specific resources on the File Server.
IIS (Web Server)	Also called an Application Server, IIS serves up web pages and provides you the ability to host multiple web sites along with running web-based applications. IIS 6.0, which comes with Windows Server 2003, is dramatically more secure than previous versions. IIS is no longer installed by default and is locked down by default when installed on a server.

Deploy the security configuration for servers that are assigned specific roles

Once you have identified the server roles that will be used in your network, you then need to configure security. Microsoft has made this easy through the use of **Group Policy** and **Security Templates**. Security Templates allow you to use a pre-defined Security Template, customize it for your desired security settings and then apply the settings through Group Policy. You can link the Group Policy that has the Security Template configured to the local computer itself, the Active Directory site it belongs to, the domain or, generally the best option, to an **Organizational Unit (OU)**. This way if you design your OU structure correctly, your task of applying security configuration to servers will be much easier to deploy as well as troubleshoot.

Create custom security templates based on server roles

With Windows Server 2003, Microsoft includes several default security templates. Security templates are pre-defined, built-in security policies. Templates can be applied to your local computer or to Group Policy objects. For example, there are templates to increase security, provide backwards compatibility for older applications and restore a default security configuration, to name a few. The great thing about having these default security templates is that you can use a pre-defined template as a starting point and then customize it by adding your own settings. Once you have made all the changes to properly secure your servers, you save the template as a different name. Then you have to apply that template either via Group Policy or the **Security Configuration and Analysis** utility to link it to a server object. Below is a list of the default security templates that you will be able to build a custom template from.

Template	Description
Default security <u>Setup security.inf</u>	Default computer security settings that were applied during the installation. Clients and member servers can use this template, but not domain controllers (DCs).
Domain Controller default security <u>DC security.inf</u>	Created during the installation of Active Directory (AD), when the server is promoted to a DC. It contains the DCs' Registry, file, and system service default security settings.
Compatible <u>Compatws.inf</u>	Relaxes default file and Registry permissions to ensure maximum application compatibility for the Domain User's group.
Secure <u>Secure*.inf</u>	Medium-level security settings with minimal impact to application compatibility. Enhances security settings that are least likely to impact application compatibility. Best used in environments containing no down-level clients.
Highly Secure <u>hisecc*.inf</u>	Highly secure template that require strong encryption and signing for a secure channel. To use highly secure settings, domain controllers must use Windows 2000 or Windows Server 2003. Increases in security settings are likely to affect older clients and applications from working properly.
System root security <u>Rootsec.inf</u>	Specifies permissions for the root of the system drive. Use this template to reapply and restore default root directory settings or to apply setting to other disk volumes.

To customize a predefined security template, perform the following:

1. In a Microsoft Management Console (**MMC**), add the **Security Templates** snap-in.
2. In the console tree, expand Security Templates, and then double-click the default path folder (systemroot\Security\Templates).
3. In the details pane, right-click the predefined template you want to modify, and then click **Save As**.
4. In the Save As dialog box, type a new file name for the security template, and then click Save.
5. In the console tree, double-click the new security template to display the security policies, and navigate until the security attribute you want to modify appears in the details pane.
6. In the details pane, right-click the security attribute, and then click **Properties**.
7. In the Properties dialog box, select the **Define this policy setting in the template** check box, make your changes, and then click OK.
8. In the console tree, right-click the new security template, and then click **Save**.

To create a new security template and install the Security Configuration and Analysis MMC snap-in, do the following:

1. In an MMC console, add the Security Templates snap-in.
2. In the Add/Remove Snap-in dialog box, also select Security Configuration and Analysis and then click the OK button.
3. In the console tree, expand Security Templates, right-click the default path folder (systemroot\Security\Templates), and then click **New Template**.
4. In the Template Name box of the systemroot\security\templates dialog box, type the template name and an optional description and then click OK.
5. In the console tree, double-click the new security template to display the security policies, and navigate until the security attribute you want to modify appears in the details pane.
6. In the details pane, right-click the security attribute, and then click **Properties**.
7. In the Properties dialog box, select the **Define this policy setting in the template** check box, make your changes, and then click OK.
8. In the console tree, right-click the new security template, and then click **Save**.

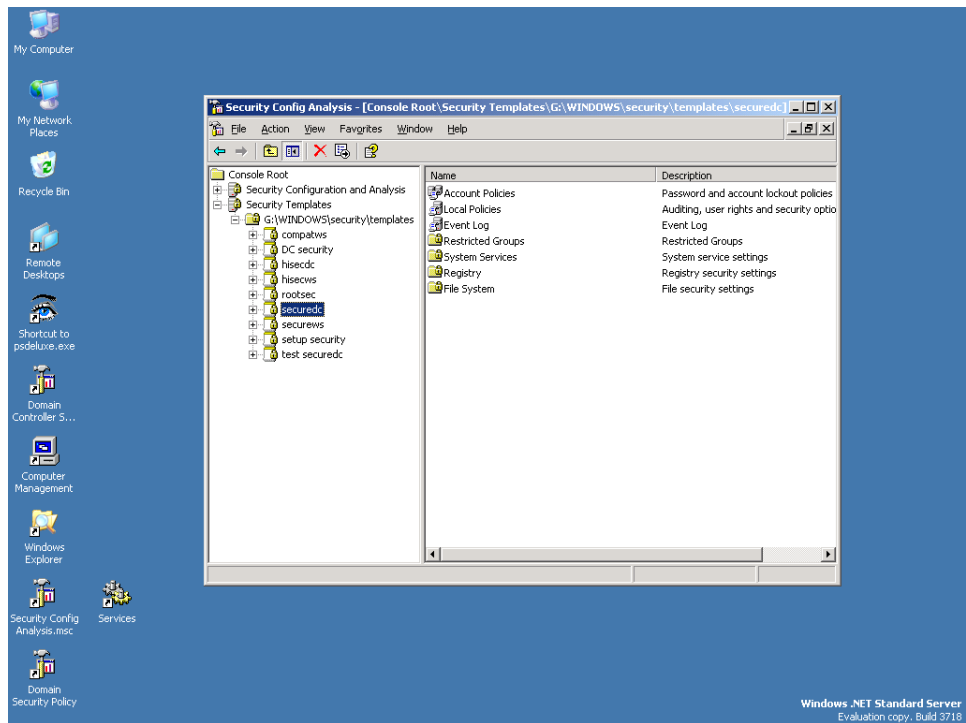


Figure 3 - The Security Configuration and Analysis MMC snap-in

Figure 3 depicts the **Security Templates** and the **Security Configuration and Analysis** MMC snap-in displaying the *securedc* template, saved as *test securedc*.

Evaluate and select the operating system to install on computers in an enterprise

Remember that the newer the Microsoft client operating system the more security features it is going to contain. Because of this, you should only select operating systems that make it easier for you as the administrator to configure security. These types of operating systems would be the Windows 2000 family, Windows XP Professional, and the Windows Server 2003 family. All of these clients systems allow you to apply Group Policy and operate very well in an Active Directory enterprise Domain.

To analyze your server system security settings:

1. In the security console, right-click **Security Configuration and Analysis** and choose **Analyze Computer Now**.
2. In the Perform Analysis dialog box, accept the default location for the error log file path and click the OK button to continue.
3. When the analysis completes, expand all nodes in the left pane under Security Configuration and Analysis. In the right pane, view and compare your security database policy settings with the current settings on your local Windows Server 2003.

Note that when you analyze security settings, no changes are made to your system. The results of the security analysis show the differences in your custom security template as compared to your actual computer system settings.

Identify the minimum configuration to satisfy security requirements

Each organization is going to have different security requirements. You will want to listen to several key groups of individuals before deciding on what the minimum amount of security is going to be. These groups of people should include individuals from these departments:

- IT
- Managers
- HR
- Help Desk
- Executive
- Users from different departments

You will need to weigh the pros and cons of the security that needs to be applied and funnel out any of the unnecessary wants that may prohibit users from completing their work.

When security recommendations are completed, it's time to analyze selected client computers' security settings and make appropriate adjustments.

Planning, Implementing, and Maintaining a Network Infrastructure

Plan a TCP/IP network infrastructure strategy

A Windows Server 2003 Network Infrastructure can appear very complex at first. However, once you learn what the core services are and why they are needed, it will be much easier for you to pass the exam associated with this exam manual, as well as deal with issues you'll encounter with TCP/IP network infrastructures.

As a system engineer or administrator your time will be split among many tasks including planning, implementing, and maintaining your network infrastructure. Out of those three phases you can expect to spend about 20% time planning, 15% implementing and 65% maintaining your network infrastructure. Even though there is a smaller amount of time allocated for planning and implementing, if you don't do these tasks correctly the first time, you will waste countless hours later troubleshooting attempting to fix TCP/IP issues.

TCP/IP is the protocol of choice these days. It is the de-facto protocol of the Internet. Without it, the Internet could not exist, nor could a Windows Server 2003 Network. With so many of the services of Windows Server 2003 being dependent upon TCP/IP, Windows Server 2003 installs this protocol by default. In addition to TCP/IP, you can also install other protocols for compatibility with other operating systems, networks, and other services as needed.

A Windows Server 2003 computer needs two items before it can properly communicate on a TCP/IP network: **IP address** and a **subnet mask**. If you want to communicate with other networks, then you will also need to configure a **default gateway**.

Configuring TCP/IP

By default, TCP/IP is installed when you install Windows Server 2003.

There are two ways to configure TCP/IP: using a **DHCP** server to automatically assign IP addresses to clients; and **manually** configuring TCP/IP settings. This latter option is useful for various servers that run services such as WINS, DNS, and DHCP requiring a static IP address.

To configure TCP/IP, you can run **Network and Dial-up Connections**, double-click **Local Area Connection**, select **Properties**, choose the **General** tab and then click **Properties** for the **Internet Protocol (TCP/IP)** item. This dialog box allows you to either automatically receive or manually set these items:

- IP address
- Subnet mask
- Default gateway
- Preferred and alternate DNS servers

The **Advanced...** button provides access to these items:

- IP addresses
- Default gateways and interface metric
- DNS suffix and registration settings
- WINS settings
- IP Security and TCP/IP filtering options

As mentioned above, every systems engineer or administrator must decide whether he or she wants to manually assign IP addresses for the clients or use a service like DHCP to dynamically configure the IP addresses of the TCP/IP hosts within the network.

If you have more than 5 clients, I would recommend that you dynamically configure the clients IP settings. However, for servers you will commonly assign a static IP address to ensure that the IP address will not change, a stipulation for many network services.

DHCP

If you have more than 5 clients and want to configure client IP settings dynamically, then your best bet is to use Windows Server 2003's DHCP service. DHCP stands for **Dynamic Host Configuration Protocol**. This protocol was developed to ease the headaches involved with manually managing client IP addresses, such as duplicate IP addresses and other TCP/IP setting inaccuracies. DHCP is based on the **BOOTP** protocol. DHCP enables dynamically distributing IP addresses, and other associated configuration data, through an open standard defined by RFC 2131 and RFC 2132.

Because DHCP has its roots in the BOOTP protocol, DHCP communications do not traverse routers by default. To traverse routers either **BOOTP forwarding** has to be set on the routers or **DHCP Relay Agents** must be set up on the segments without a DHCP server.

Analyze IP addressing requirements

Most networks already have an IP addressing scheme in place. It is very rare for a systems engineer to go into an existing company and not have a functional IP based network. However, you may have to reevaluate an IP based addressing scheme to optimize network performance and fault tolerance.

Your network likely consists of different types of devices from XP workstations, NT, 2000 and 2003 servers, routers, switches and printers to name just a few. Each will be identified as a TCP/IP host and have a unique logical IP address assigned to them.

You can specify IP addresses based on:

- **Classes** (A, B, C) with an associated default mask
- Classes with variable length subnet masks (**VLSM**)
- Classless Inter-Domain Routing (**CIDR**) with a specified prefix length

The following table lists the class-based TCP/IP addresses ranges, default subnet mask and purpose.

Address Class	Address Range	Default Subnet Mask	Purpose
A	1-126.xxx.xxx.xxx	255.000.000.000	Host/Network
B	128-191.xxx.xxx.xxx	255.255.000.000	Host/Network
C	192-223.xxx.xxx.xxx	255.255.255.000	Host/Network
D	224-239.xxx.xxx.xxx	None	Multicast
E	240-255.xxx.xxx.xxx	None	Experimental

When you are analyzing your network, you will probably discover that you need to separate your private network from the public Internet. One big reason why you want to do this is for security. If you had all your TCP/IP hosts connected with a public IP address, it would be like giving a hacker the code to your ATM card, not to mention that it would be extremely expensive to purchase that many public IP addresses.

There are private ranges of IP addresses that can be used securely and freely in your private network. These private IP addresses are used for internal networks and are non-routable. This gives companies an almost endless supply of valid IP addresses within the private network. Later we will discuss how these TCP/IP hosts can connect to the Internet while using a private IP Address. The table below lists the private ranges defined in RFC 1918.

Address Class	Address Range
A	10.0.0.0 - 10.255.255.255
B	172.16.0.0 - 172.31.255.255
C	192.168.0.0 - 192.168.255.255

Plan an IP routing solution

IP Routing in a Windows Server 2003 Network Infrastructure is something you are definitely going to want to spend some time on. Microsoft expects system engineers to be able to plan an IP routing strategy for a company that consists of multiple networks. Be prepared to know how to plan a routing solution from a corporate office to multiple branch offices. Also, make sure you know about the proper networking devices and where they belong according to the OSI model. Listed in the table below are network devices, the OSI (Open system Interconnect) layer they communicate in and their definition.

Device	OSI Layer	Definition
Hub	Physical Layer	Does not Process data. Extends the network by retransmitting the signal.
Switch	Data-link Layer	Forwards frames according to the destination address. Uses temporary or virtual connections to connect source and destination ports.
Layer 3 Switch	Network Layer	Is a limited-purpose hardware-based IP router with bridging capabilities. Also Performs layer 2 switching.
Router	Network Layer	Used to link WANs and dissimilar LANs. Operates at the packet level. Sends packets based on packet addressing.

Traditionally routers were used at the outer edges of networks. This is no longer the case and you will find many larger networks broken up in to smaller networks by the use of routers. Windows Server 2003 routing is used at the IP level since it is the protocol that is used by most network infrastructure services. Routers also isolate networks from each other and are the foundation for creating **Screened Subnets**. A screened subnet consists of the portion of the network that resides between the public Internet and private network. However, routers alone will not secure your network from malicious attacks. Your security plan needs to include firewalls, filtering and **intrusion detection software** (IDS).

When planning an IP routing solution it is important to understand Collision domains and Broadcast domains.

- Collision domain
 - ▶ A collision occurs when two computers on the same Ethernet network node try to send data at the same time. As the network can only carry data from one computer at a time, a collision occurs. When this happens, both computers back off and wait a random period and then retransmit. The random period will only be a few milliseconds but enough for one computer to get its data in while the other waits. A collision domain is an area of the network where a particular group of computers are likely to have collisions.
 - ▶ Collisions do not traverse a bridge or a switch, so a switch forms an edge or border of a collision domain.
- Broadcast domain
 - ▶ A broadcast domain is a logical area in a computer network where any computer connected to the computer network can directly transmit to any other in the domain without having to go through a routing device. More specifically, it is an area of the computer network made up of all the computers and networking devices able to communicate together by sending a frame to the data link layer broadcast address.

Routers are generally more expensive than switches, and, as such, can become expensive when managing broadcast domains. **Virtual LANs** (VLANs) can provide you with a cheaper solution because they allow switches to also contain broadcast traffic. A VLAN is a LAN that the broadcast domain creates artificially through the switch configuration. Clients don't need to be connected to the same switch or even be located in the same physical area because the networking devices on your network will represent these VLANs as actual LANs. Realize that not all switches can do this.

Create an IP subnet scheme

A **subnet mask** is a method used in TCP/IP to divide the network portion of an IP address from the host portion. A subnet is a portion of a TCP/IP network that accesses other subnets via routers. A shorthand way to describe a subnet mask is to append a forward slash ("/") after the IP address followed by a number that indicates how many bits of the IP address belong to the network portion. For example, the address 120.1.2.3/8 represents an IP address of 120.1.2.3 that has a subnet mask of 255.0.0.0. The "/8" is read as "slash 8" or "eights." This method of naming is called **classless addressing**.

Legacy networking hardware and software prevent the use of all 0's or all 1's in a subnet. With recent hardware, though, whether or not all 0 or 1 subnets are allowed depends on the routing protocol in use.

CIDR and VLSM

Classless Inter-Domain Routing is a method that allows the use of a single IP address and subnet to reference a group of addresses. The address is called "Classless" because it ignores the "Classful" conventions that employ three classes of addresses: Class A, B and C, which state that the first 8 bits of a Class A address reference the network, the first 16 bits of a Class B address reference the network and the first 24 bits of a Class C address reference the network.

With CIDR, between 13 and 27 bits of an IP address can refer to the network segment while the remaining bits refer to the host portion. In the case of 27 bits, this would result in a network segment with 30 hosts ($2^{32-27} - 2$).

These are the major benefits to CIDR:

- Fewer entries are needed in routing tables since several **contiguous** network segments can be “rolled up” or aggregated into one IP address.
- CIDR makes more IP addresses available. By ignoring the rigid rules of “classful” addressing, different combinations of network segment / host number pairings are possible.

A CIDR address uses notation that takes the form of **w.x.y.z/s**, as in 192.168.8.0/20 where the “/20” indicates the number of bits that mask off the “network” portion of the address. Note in this example that the subnet mask is shorter than the one used for the standard “classful” address.

CIDR is used primarily to optimize routing tables. For example, let’s assume that four **contiguous** Class C networks, 192.168.8.0/24 to 192.168.11.0/24, are reachable via a single router. Instead of having 4 routing table entries to describe the route to these networks, we can use one routing table entry to specify an interface to reach the networks described by a single summarized entry of 192.168.8.0/22. In large, enterprise-wide configurations, proper planning of IP address space to allow the use of CIDR blocks can result in significant savings of resources used to maintain routing tables.

Variable Length Subnet Mask helps large organizations to better allocate the IP address space. VLSM, which is closely related to CIDR, allows various network subnets to have different subnet masks. By properly designing your IP address space, subnet masks with fewer bits assigned to a network address can be used to summarize or aggregate subnet masks with more bits assigned to the network address. This results in routers having smaller route tables.

In summary, CIDR allows routers to group routes together to reduce the quantity of routing information carried by the core routers, whereas VLSM helps optimize available address space.

Plan and modify a network topology

When planning or making modifications to a network topology, the first thing any good administrator needs to do is to document what happens. Even after the network is set up, you should still be documenting day-to-day processes to ensure that these processes map to the initial design, as well as to create a record of the latest network configuration.

There is not just one plan that you will use when you plan a network topology. Instead, you will use several plans, including the ones listed below.

Type of Plan	Brief Description
Capacity Plan	Outlines a course of action to ensure that the solution performs in an acceptable manner to the users
Deployment Plan	Documents the entire deployment strategy, any contingencies, and support tools required
Security Plan	Ensures access to data, services, and resources is constrained according to the project's guidelines
Pilot Plan	A select group of users that will participate in a small pilot test run of the new network
Test Plan	Outlines the strategy that will be used to test the solution

During all planning phases you must identify key areas of your network. These key areas may contain specific network resources or different protocols that will be used.

Plan the physical placement of network resources

Most companies physically secure their servers in a locked room that only administrators have access too. This is important because you don't want anybody to go up to one of your servers and start eavesdropping, or worse, turn the servers off by unplugging them or damaging them. Servers are NOT the only devices in your network that you need to plan for physical security. You will also want to lock up your network switches and routers. If a user or hacker achieves physical control of these devices then you lose control of them.

When planning the physical placement of your network resources consider the different technologies that are available. Listed below are a few tools to consider:

- Locks
- Smart Card entry
- Retinal scans
- Finger print scans
- Voice print readers

Any one of these methods makes it a lot more difficult to break into the room that contains the protected devices. You may also want to consider multiple locations for stationing of network resources in case there is an emergency.

Identify network protocols to be used

Windows Server 2003 comes installed with TCP/IP. TCP/IP is required for an Active Directory environment. Depending on the other types of clients and servers, you may also have to use other protocols. TCP/IP will be discussed further throughout this Exam Manual; however, if you have Novell clients and servers, you will also need to know how to set up additional protocols to communicate effectively between your Novell and Microsoft networks.

Windows Server 2003 and Novell NetWare integration uses the **NWLink** protocol for **IPX/SPX**-compatibility. This allows Windows Server 2003 systems to communicate with older NetWare servers that do not use TCP/IP. Because older versions of Novell Netware use the **NCP** protocol while Windows uses the **SMB/CIFS** protocol, these two platforms are unable to share files without special configuration. Although a Windows Server 2003 system can communicate with a Netware server running a client/server application, and a Netware server can communicate with a Windows client/server application, compatibility issues crop up when dealing with file and print services.

In addition to NWLink, these other services enhance connectivity with Netware:

- **Client Services for NetWare (CSNW)** – Enables Windows clients to connect directly to NetWare shares on Novell servers, without needing a native Netware client.
- **Gateway Services for NetWare (GSNW)** – Enables Windows clients, without needing a native Netware client, to connect to Netware shares via a gateway set up on a Windows Server 2003. If several clients use this method to access NetWare data shares, consider using CSNW or adding the native NetWare client, Client 32. Otherwise, the GSNW may become a bottleneck.
- **File and Print Services for NetWare (FPNW)** – Allows NetWare client computers to access Windows Server 2003 network shares. In addition, the **Microsoft Directory Synchronization Services** (including the **Directory Service Manager for Netware** and the **Directory Service Migration Utility**) and the **File Migration Utility (FMU)** help to synchronize **Active Directory** with **NDS** (Novell Directory Services) as well as migrate from NDS to Active Directory, and migrate a Netware file system to Windows Server 2003.

No matter which protocol you decide to use, you will need to make sure that you are using the preferred protocol first. Network Bindings represent the **order** in which protocols are bound to the NIC (Network Interface Card) as clients and servers attempt to communicate. Network communication is attempted in the order of the network protocols in the binding order until a common protocol is found between both the client and server.

To view or change the binding order of your protocols:

1. Right-click My Network Places and choose Properties.
2. On the Network Connections menu bar select Advanced and then click Advanced settings...
3. On the Adaptors and Bindings tab, in the Bindings for Local Area Connection, text box, select the protocol you want to adjust and then click the up or down button accordingly.

There are two types of objects, or providers that relate to bindings:

1. **Network** – For example, Microsoft Windows Network, Gateway Services for NetWare
2. **Printer** – For example, LanMan Print Services, HTTP Print Services

For optimization purposes, keep these recommendations in mind:

- The binding order should be from the most often used protocol to the least, so that a common language can quickly be found
- Unneeded protocols should be removed to reduce needless network traffic

Plan an Internet connectivity strategy

There are several solutions that a company can use to provide Internet connectivity. Each company will have to analyze current security policies, network access and budgetary constraints when deciding what type of Internet connectivity they will choose.

When providing Internet connectivity from your private network, the first thing you need to address is security. By opening a connection to the Internet you are exposed to allowing potentially harmful data into your network. It is not only your users that are at risk. You may inadvertently expose a network service, such as DNS, to the Internet when it shouldn't be. This could give potential hackers a virtual blueprint of your network. One of the easiest ways to protect your network is to use some form of a firewall for protection.

Every company should have what is known as an *acceptable internet usage policy*. This policy will define what users can and cannot do on the Internet. It should be reviewed by management, HR, IT and then signed by the employees.

Windows Server 2003 simplifies Internet access by including several technologies that can be used to securely connect to the Internet.

Connecting to the Internet by using a **Router**:

- For smaller networks, Windows Server 2003 can function as a router running RRAS (**Routing and Remote Access**)
- Routers read packet headers and select the best path available
- Routers enable all users on a network to share a single connection to the Internet
- A Router is the easiest method for connecting a private network to the Internet, but realize that it does not secure your network against unauthorized access

Connecting to the Internet by using a **Firewall**:

- Windows Server 2003 comes installed with a limited firewall
- Firewalls can be software or hardware based
- By default, firewalls block all incoming and outgoing traffic
- Firewalls can be very difficult to configure

Connecting to the Internet by using Network Address Translation (**NAT**):

- For smaller networks, Windows Server 2003 running RRAS can function as a NAT Server
- NAT reduces IP address registration cost, conceals internal IP addresses from external networks, and provides limited DHCP and DNS features

Connecting to the Internet by using Internet Connection Sharing (**ICS**):

- Windows Server 2003 can function as an ICS Server
- ICS is great for very small offices or home offices, but cannot be used in a network that has DNS, DHCP, or Active Directory

Connecting to the Internet by using a **Proxy Server**:

- Windows Server 2003 does not have a built in Proxy server
- Microsoft does have Internet Security and Acceleration (ISA) server. This can:
 - Cache content for faster access to previously used sites
 - Block potentially unsafe websites
 - Provide great logging features

Plan network traffic monitoring. Tools might include Network Monitor and System Monitor

Two of the largest contributors to network traffic are the *number of hosts* on the network and the *number of services* provided by network servers. Other culprits are *protocols* and *network shares*. Thus, the easiest way to reduce traffic is to remove any of the above items that are not necessary. Also, consider segmenting your network.

Keep in mind that the **Network Monitor** product that comes bundled with Windows Server 2003 is a limited edition that captures frames only between the host and other computers, not all frames present on the network segment. The full version of Network Monitor, which overcomes this limitation, ships with Microsoft's SMS product.

Network Monitor has two components that enable it to capture and analyze data packets:

- **Network Monitor** – The GUI tool that captures and analyzes frames and data packets
- **Network Monitor Driver** – This is installed, along with Network Monitor, on a server. You may want to install just this driver on a client so that the SMS version of Network Monitor can monitor the client

You can install Network Monitor by following these steps:

Select in the control Panel **Add/Remove Windows Components -> Management and Monitoring Tools -> Network Monitor Tools.**

You can install the Network Monitor driver by following these steps:

Select **Network and Dial-up Connections -> Properties for the Local Area Connection -> Install -> Protocol -> Add -> Network Monitor Driver.**

System Monitor, formally called Performance Monitor in Windows NT, is used to monitor various performance objects like the CPU, memory and hard drive by selecting various counters and charting or reporting the results using Performance Logs and Alerts.

System Monitor is installed by default and is accessed by selecting **Performance** from the Administration Tools menu. By default, System Monitor displays the Page/sec, Avg. Disk Queue Length and % Processor Time counters in a graph format as shown in *Figure 4*. You can create and add new counters by clicking the plus (+) toolbar button. Because there are many different counters for each performance object, the Add Counters dialog box includes an Explain button that details the features of the specific counter that you selected.

Counter logs allow you to create, change, view or delete your saved results from System Monitor. A sample binary file is included and displays on startup. **Trace logs** and **Alerts** can be created and when a defined threshold is reached, can inform the administrator via e-mail.

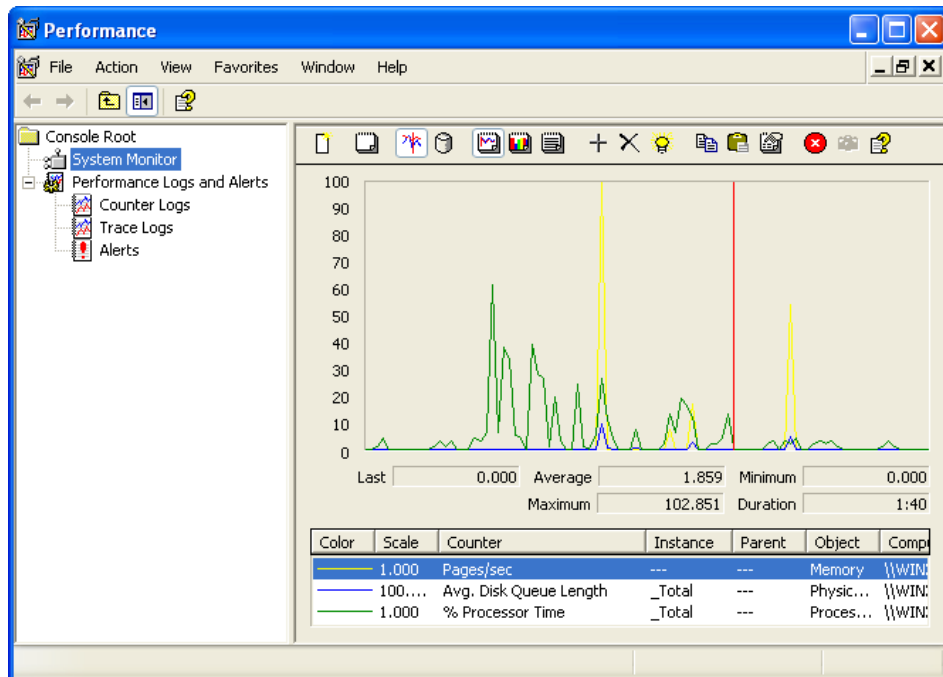


Figure 4 - System Monitor, Performance Logs and Alerts

Capturing Data

Refer to this Microsoft Knowledge Base article for details on how to capture network traffic using Network Monitor. (KB# [Q148942](#))

Being able to capture network data, analyze the results, and optionally save the results to a file is important. However, of even more importance is creating a **baseline** of captured data on a regular basis. It is the comparison of a network data capture when compared to an earlier baseline that will likely help you most in troubleshooting a problem.

Troubleshoot connectivity to the Internet

At times it seems as if they're as many troubleshooting methods as there are troubleshooters. The connection type you have to the Internet will greatly determine the method you will use to troubleshoot your connectivity to it. The first step is to identify that connectivity method.

Diagnose and resolve issues related to Network Address Translation (NAT)

As its name describes, Network Address Translation or NAT translates IP addresses for hosts on the local private network so they can communicate on the Internet. There are three **pools of private addresses** that are used by NAT:

- 10.0.0.0 to 10.255.255.255
- 172.16.0.0 to 172.31.255.255
- 192.168.0.0 to 192.168.255.255

When diagnosing problems with NAT make sure you check several things, including:

- Ensure that a DHCP Scope is not conflicting with the client addresses NAT is assigning.
- Ensure that you have the correct interfaces assigned to the Public and Private interfaces in NAT. You can review this setting in RRAS, under IP Routing.
- If applications are not working properly, check your Mappings to ensure they are using the correct ports.
- If you have multiple public IP addresses ensure that the **translation of tcp/udp headers** items are checked.

Diagnose and resolve issues related to name resolution cache information

Name resolution problems can occur if a client computer is resolving names incorrectly. Your computer has a local name cache that it will check before using any other name resolution method. To view your host name cache, do the following:

1. Open the command prompt
2. Type `ipconfig /displaydns`

This will show your current local name cache for host names. To clear your host name resolution cache, you must do the following:

1. Open the command prompt
2. Type `ipconfig /flushdns`

This will flush your current host name cache and, if there were any problems, they should be fixed as long as your DNS server or HOSTS file is not configured incorrectly.

If you want to view your NetBIOS name cache then follow these steps:

1. Open the command prompt
2. Type `nbtstat -c`

This will display your NetBIOS name cache. To clear your NetBIOS name cache and hopefully fix a NetBIOS name cache problem, follow these steps:

1. Open the command prompt
2. Type `nbtstat -RR`

Diagnose and resolve issues related to client configuration

If your clients are having problems related to the TCP/IP configuration then it is probably one of two things. First, if you manually configured the client's settings you may have mistyped a number in the IP settings. Otherwise, you may have misconfigured your DHCP server to give out incorrect settings that are now being configured on your clients. To view your current TCP/IP settings on a client follow these steps:

1. Open the command prompt
2. Type `ipconfig /all`

Verify the output as to what should be configured: changes may be necessary. Depending on how your clients get their TCP/IP settings, you will have to either manually fix the TCP/IP problem on the client, or make adjustments through the DHCP manager console.

Troubleshoot TCP/IP addressing

When problems arise with TCP/IP configuration settings you will want to make sure that the clients first are able to contact a DHCP server. Next you will want to ensure that they are getting the right settings from the DHCP server. As you will learn later, DHCP scope options work in the order that they are applied therefore you may have to map out what settings are being applied where and see if that is truly the resultant DHCP settings you want it to be. Always remember that anything that is statically input will override any automatic DHCP settings.

Diagnose and resolve issues related to DHCP server address assignment

One of the biggest issues related to DHCP is miss-configured scopes. DHCP Scopes define what TCP/IP settings are going to be applied and configures those settings. For example, if you configure the Gateway address to be on a different network than what your DHCP clients are on, then they will never be able to communicate outside of their network. Make sure to look closely at every setting that you configure at the scope level. Also, it doesn't hurt to get another set of eyes on it to help you check settings.

Some of the most common scope options to examine are these:

- The IP address of a router. To issue this information, configure the **003 Router** option with the IP address of a default router. This router is commonly referred to as the default gateway address.
- The IP address of one or more DNS name servers available to clients. To issue this information, configure the **006 DNS Servers** option with the IP address of one or more DNS servers.
- The DNS domain name. A DNS domain name defines the domain to which a client computer belongs. The client computer uses this information to update a DNS server so that other computers can locate the client. To issue this information, configure the **015 DNS Domain Name** option with the proper DNS domain name.
- The IP address of one or more WINS servers available to down level clients. The down level client uses a WINS server for **Network Basic Input/Output System (NetBIOS)** name resolution. To issue this parameter, configure the **044 WINS/NBNS Servers** option with the IP address of one or more WINS servers.
- The type of NetBIOS over TCP/IP name resolution. To issue this information, configure the **046 WINS/NBT node type** option with the appropriate NetBIOS node type. The type of name resolution determines the order in which a client uses NetBIOS name servers and broadcasts to resolve NetBIOS names to IP addresses. 0x1, the default issues broadcasts and should be changed to 0x8 node type to query the WINS server first for name resolution.

DHCP scopes also work in an order of precedence. This means just because you may have set a scope option for the entire server to use the 006 DNS server of 192.168.1.10, if there is an option on a lower being configured, then the option applied last will be applied. This precedence order is as follows:

- **Server-level** options apply to all DHCP clients that lease an IP address from the DHCP server.
- **Scope-level** options are available only to clients that lease an address from that scope.
- Options that you configure at the class level are available only to clients that identify themselves to the DHCP server as belonging to a particular **Class-Level**.
- Options that you configure at the **Reserved-Client-Level** apply to specific clients.

Any option that is manually configured will always override settings that a DHCP server has configured.

Plan a host name resolution strategy

Ever since Windows 2000, DNS host name resolution has been the preferred method of resolving names for Microsoft clients and servers. There are a couple ways to implement host name resolution into your network.

The first would be to use a static **HOSTS** file. This is not the preferred method since it involves a lot of administrative work. The HOSTS file is a static file that is stored in the %systemroot%/System32/Drivers/Etc directory. This file would have to be manually updated each time a new host was added to the network and this change would then have to be applied to each client and server. As you can see, it is way too time consuming to use in a network of most any size.

The second method is to use the **Domain Name System (DNS)**. This service has the ability to dynamically register clients and servers into its database so that the administrator does not need to do this manually. The size and services that you may be running in your network will ultimately decide which of the two

methods of host name resolution you will use.

Plan a DNS namespace design

The DNS server can store and use many different types of records. They are listed in the table below:

Record Type	Purpose
A	Host address record – for mapping a DNS name to an IP address. Requires name, IP address and whether or not to “Create associated pointer (PTR)” record.
CNAME	Canonical Name - an alias domain name for a name already specified as another resource type in the zone. An example is to assign an alias name “www2” to an FQDN for target host “eng.brainbuzz.com”.
MB	Mailbox record.
MG	Mail group record.
MINFO	Mailbox or mailing list information - usually used to specify a mailbox for error messages.
MX	Mail exchanger record - details message routing to a mail exchange host. Entry fields include the domain name, the name of the mail server and the priority which defaults to 10 (the lower the priority number, the higher the priority of this particular mail server).
NS	Name server record - specifies a server that is authoritative for a certain zone.
PTR	Pointer record - used in reverse lookups. Remember that the DNS installation does not create a reverse lookup zone by default.
TXT	Text record - can hold descriptive text that can be applied to a specific DNS name.
RT	Route Through - details intermediate-route-through binding for hosts that do not have their own WAN address.
SRV	Service Record - used by Windows 2000 for Active Directory and “Dynamic DNS”. Active Directory can work with non-Windows 2000 DNS servers so long as those DNS servers support the use of SRV records.

Plan zone replication requirements

There are four replication options for **Active Directory-integrated DNS zones**. These can be selected when the zone is created or when the administrator wants to change the storage method for an existing zone. When deciding which replication option to choose, consider that the broader the replication scope, the greater the network traffic caused by replication. For example, if the administrator chose to have Active Directory-integrated DNS zone data replicated to all DNS servers in the forest, this will produce greater network traffic than replicating the DNS zone data to all DNS servers in a single Active Directory domain in that forest.

Replication options for Active Directory-integrated DNS zones are listed as follows:

Replication Option	Description
All DNS servers in the Active Directory forest	The zone data is replicated to all the DNS servers running on domain controllers in all domains of the Active Directory forest.
All DNS servers in a specified Active Directory domain	The zone data is replicated to all DNS servers running on domain controllers in the specified Active Directory domain. This option is the default setting for Active Directory-integrated DNS zone replication.
All domain controllers in the Active Directory domain	The zone data is replicated to all domain controllers in the specified Active Directory domain regardless of whether there are DNS servers running on the domain controllers in the domain. If the administrator wants to have Windows 2000 Server DNS servers load an Active Directory-integrated DNS zone, this replication scope setting must be selected for that DNS zone.
All domain controllers specified in the replication scope of an application directory partition	The zone data is replicated to all the domain controllers specified in the replication scope of the application directory partition.

Given any scenario, one might be able to argue that replicating DNS zone data to every DNS server in the forest might cause excessive traffic; however, the trade off is extreme ease of administration and deployment. Especially in smaller organizations, where an IT department may not exist, configuring DNS for Active Directory can be difficult for most. In these situations, the best solution should be the easiest solution: using forest-wide replication for any DNS zone is considered the best option. If every zone is replicated to all DNS servers in the forest, then a customer should not need to use stub zones, secondary zones, or forwarders when configuring internal DNS resolution using Active Directory.

Figure 5 depicts the New Zone wizard, AD Zone Replication Scope dialog box with the default option selected.

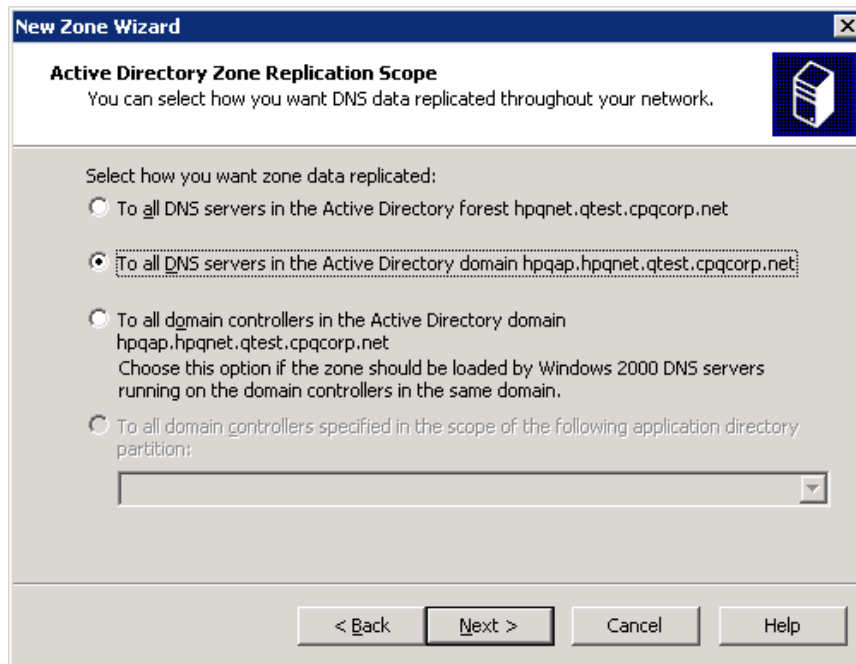


Figure 5 – New Zone Wizard

Plan a forwarding configuration

In Windows Server 2003, Microsoft has improved upon the forwarding of Windows 2000. The new forwarding feature added in Windows Server 2003 is called **Conditional Forwarding** and allows a DNS server to forward queries to other DNS servers based on the DNS domain names in the queries. With conditional forwarding, a DNS server can be configured to forward all the queries it receives for names ending with `widgets.microsoft.com` to a specific DNS server's IP address, or to the IP addresses of multiple DNS servers.

For example, when two companies (`example1.com` and `example2.com`) merge or simply collaborate, they may want to allow clients from the internal namespace of one company to resolve the names of the clients from the internal namespace of another company. The administrators from one organization, such as `example1.com` may inform the administrators of the other organization, such as `example2.com`, about the set of DNS servers that they can use to send DNS queries to for the name resolution within the internal namespace of the first organization. In this case the DNS servers within the `example2.com` organization will be configured to forward all queries for names ending with "example1.com." to the designated DNS servers.

Authoritative DNS servers cannot forward queries according to domain names for which they are authoritative. For example, the authoritative DNS server for the zone `widgets.microsoft.com` cannot forward queries according to the domain name `widgets.microsoft.com`. If the DNS server were allowed to do this, it would nullify the server's ability to respond to queries for the domain name `widgets.microsoft.com`. The DNS server authoritative for `widgets.microsoft.com` can forward queries for DNS names that end with `hr.widgets.microsoft.com`, if `hr.widgets.microsoft.com` is delegated to another DNS server.

Figure 6 shows the DNS server Properties dialog box, Forwarders Tab where DNS Domains and DNS Conditional Forwarders are added.

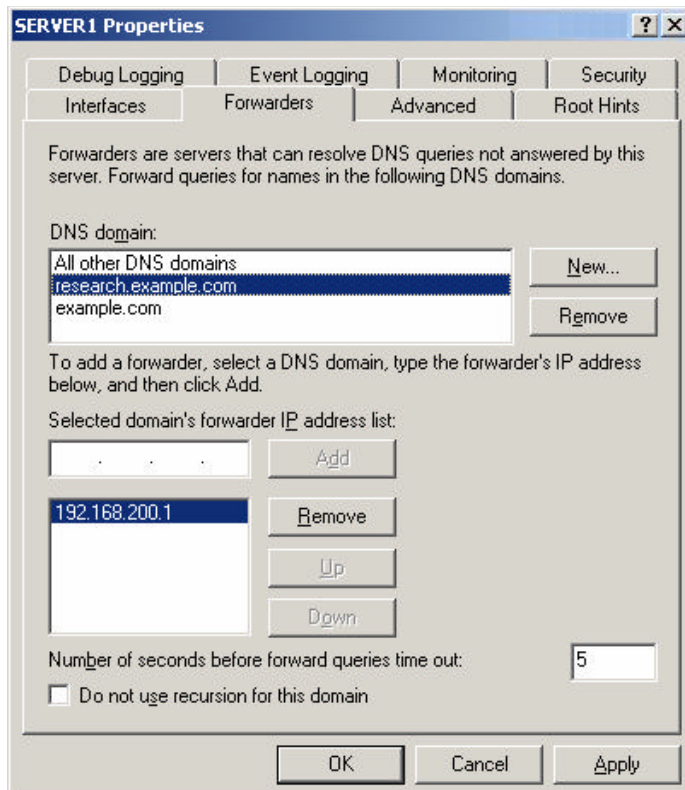


Figure 6 – Configuring DNS Forwarding

A DNS server can also be configured to not perform recursion after the forwarders fail by checking the **Do not use recursion for this domain** check box in *Figure 5*. If it does not get a successful query response from any of the servers configured as forwarders, then it sends a negative response to the DNS client. This option to prevent recursion can be set for each conditional forwarder in Windows Server 2003. For example, a DNS server can be configured to perform recursion for the domain name `research.example.com`, but not to perform recursion for the domain name `example.com`.

Plan for DNS security

Secure dynamic updates is the default method of DNS updates in both Windows 2000 and Windows Server 2003. It ensures that update requests are processed only if Active Directory authorizes them. This prevents the type of DNS attack that involves entering invalid data into the zone files. This is extremely important because, if an attacker can enter invalid data, they can disrupt your network and send your clients to rogue servers without their knowledge. They can also flood the server with a constant stream of garbage data and thereby perform a type of **DoS (Denial of Service)** attack.

Figure 7 shows the DNS primary Zone Properties box, General tab.

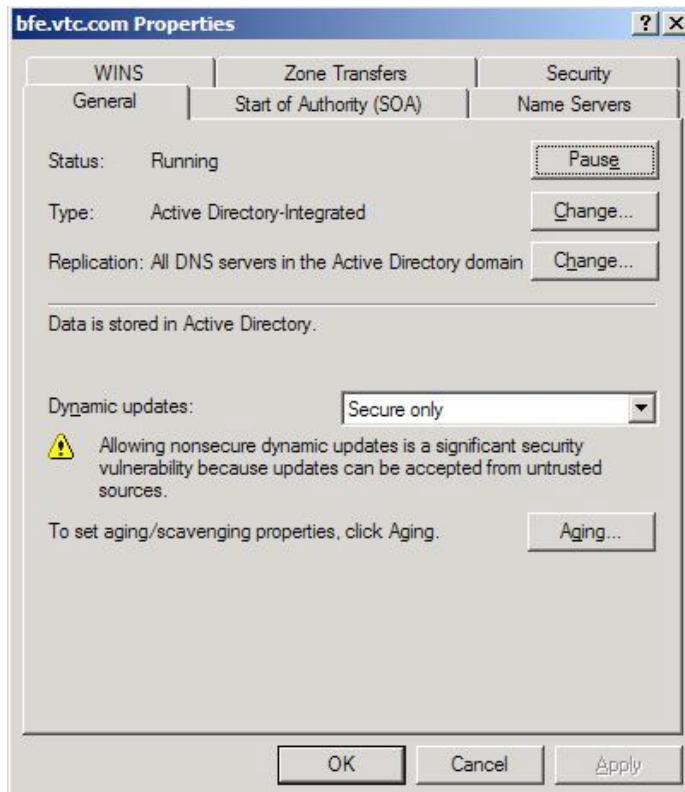


Figure 7 – Securing DNS Updates

By default, members of the Authenticated Users group have the ability to create resource records on DNS servers that are in the domain of the client computer. This enables all computers to dynamically update DNS zone data. A typical authenticated user will register a maximum of 10 records in DNS. To ensure that malicious users or applications do not create inappropriate resource records, you can set a default quota limit of 10 objects per user. This ensures that all computers can update DNS appropriately but cannot start DoS attacks.

Examine the interoperability of DNS with third-party DNS solutions

A Windows Server 2003 network does not have to rely solely on a Microsoft DNS solution. However, if you use Windows Server 2003 DNS, you can use the advanced DNS security settings that you will not get with other vendors' solutions.

A **BIND (Berkeley Internet Name Domain)** DNS platform may be able to support Active Directory depending on its version. The following table lists versions of BIND versus support for Active Directory. At minimum, BIND must support **SRV** records. These records allow clients to access services in Active Directory by mapping an IP address to the name of a server providing a service. BIND is traditionally run on **Unix/Linux** platforms but can be run on a Microsoft Server too.

DNS Server	Feature
Microsoft Windows Server 2003 DNS	<p>Integration with WINS for name resolution. Name resolution queries that fail with DNS are passed to a WINS server. (This is enabled via the Use WINS forward lookup check box on the WINS tab of the zone file properties page in the DNS.) Note that a side effect of this integration is that a WINS record in a DNS primary zone can cause a DNS zone transfer failure when that transfer is to another version of DNS, such as BIND, which isn't designed to handle WINS records.</p> <p>Secure DNS updates: the administrator can limit DNS updates to chosen secondary servers.</p> <p>Ability to integrate zones into Active Directory (Active Directory Integrated Zones). The benefit of this is that zone information management can piggyback on top of built-in Windows 2000 replication and fault tolerance. Even without Active Directory Integrated Zones, speed improvements are achieved via master servers notifying secondary ones of changes (notification-driver) as opposed to depending on polling intervals.</p> <p>A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative name server.</p>
BIND 9.2.0	Includes many new critical security updates to previous versions of BIND. It includes support for DNSSEC, IPSec and Split Zones .
BIND 8.2.1	IXFR (Incremental Zone transfers in addition to only full zone transfers.) IXFR transfers can take place from both Active Directory-integrated servers and Standard Primary servers to Standard Secondary servers.
BIND 8.1.1	Dynamic Update DNS support (the integration of DNS and DHCP facilitates, the automatic recording of client IP address / host name mappings for Windows 2000 clients).
BIND 4.9.7	SRV record support.

You can get more information on BIND and also [download it here](#).

Plan a NetBIOS name resolution strategy

It will likely be some time before we can banish WINS from most networks. In a pristine Windows Server 2003 network, with Windows 2000 and above clients, WinSock applications, DNS and Active Directory handle name resolution without help from WINS. The **Windows Internet Naming Service**, which maps NetBIOS ("computer") names to IP addresses, is required on any Windows Server 2003 network that contains down-level (anything before Windows 2000) clients and servers, or has computers running applications that are written for the NetBIOS interface, not the WinSock interface. Moreover, the Computer Browser (i.e., Network Neighborhood) service also relies on WINS. The good news, if you still require WINS in your environment, is that Microsoft has added some enhancements to the WINS product that eliminate some of the past name resolution headaches administrators suffered from when managing it.

For really small networks, the **HOSTS** files can be used in place of DNS to map hostnames to IP addresses and **LMHOSTS** files can be used in place of WINS to map NetBIOS names to IP addresses. In many cases, the broadcast nature of the NetBIOS protocol, and the use of LMHOSTS files, can probably be tolerated. However, Windows Server 2003 WINS make life easier for all but the smallest networks. If you do need to configure or view the LMHOSTS files, the sample file (LMHOSTS.SAM) resides under `\%systemroot%\system32\drivers\etc`.

There are four main components to WINS:

1. **WINS Servers** – These receive directed messages from WINS clients, or WINS proxy computers, and perform NetBIOS to IP address name resolutions. At minimum, you should have *at least two WINS servers* (for load balancing and fault tolerance), but realize that some problems can result by having an excess of WINS servers. Moreover, you should have at least two WINS servers on a network segment that is bounded by a WAN link. Microsoft recommends *one WINS server for 10,000 users plus one other for redundancy*. When your large network justifies several WINS servers, use a spoke and hub design.
2. **WINS Clients** – These use directed communication with the WINS servers as opposed to broadcasting across the network (the DHCP 0x8 node option). Clients can be Windows for Workgroups, Windows 9X, Windows NT, Windows 2000, Windows XP, and Windows Server 2003.
3. **Non-WINS Clients** – These use **broadcasts to WINS proxy computers which in turn communicate with WINS servers**. Windows 3.1, DOS and non-Windows computers like Linux and Unix fit in this category. These clients aren't left out of the WINS world thanks to WINS proxy servers.
4. **WINS Proxies** – These computers intercept broadcasts on their subnet and communicate with a WINS server on behalf of a client. This role can be filled by these systems: Windows for Workgroups, Windows 9X, Windows NT, Windows 2000 (Server and Professional), Windows XP and Windows Server 2003. Although WINS proxy servers can help a client with name resolution, they cannot perform a registration for the client.

WINS is used to resolve NetBIOS names. NetBIOS names are **16 characters** long, with the first 15 characters used to identify a unique name (for a single user or computer) or a group name (for a set of users or computers). The last character is a hex value that identifies the type of resource that is being registered. For example, these resources are created when a computer registers with WINS:

- **00h** – WorkStation
- **03h** – Messenger
- **20h** – File Server

Several other resources are possible, such as Domain Name, **1ch**.

Client computers will attempt to resolve names in four different ways. These methods are determined on how a client's node types are configured. There are four options when configuring node types:

- **B-node** uses broadcasts for name registration and resolution
- **P-node** uses a NetBIOS name server like WINS to resolve NetBIOS names
- **M-node** combines B-node and P-node, and uses B-node as the default
- **H-node** combines P-node and B-node, and uses P-node as the default

Plan a WINS replication strategy

The main benefits of having more than one WINS server on your network is that by sharing the WINS database, these servers provide fault tolerance and load balancing.

Replication Types

- **Pull Replication** – This method is **time based** and occurs on a schedule set by the administrator. This is a good choice to use for two WINS server *separated by a WAN* since the administrator can choose the best time of the day to trigger the replication. Settings include:
 - ▶ Start time of replication
 - ▶ Interval for replication (default: 30 minutes)
 - ▶ Whether or not to configure a **persistent connection** (see the next section for more details)
- **Push Replication** – This method is **event based** and occurs when a pre-determined number of WINS database changes have occurred. A version ID number tracks each change to records in the WINS database and is used as the trigger for a push replication. The push partner does not actually send data to its partner, but only reminds the partner that the partner should pull the data. Push partners should generally be connected by *fast links*. Settings include:
 - ▶ The number of increments of the version ID before initiating a replication
 - ▶ Whether or not to configure a persistent connection (see the next section for more details)
- **Immediate Replication** – An administrator can manually force a replication by right-clicking on a partner server and choosing to send an immediate trigger

Persistent Connection

This feature allows a WINS server to be permanently connected to one or more replication partners. This results in faster replication because no time is wasted at the start of each replication in order to make a connection to a replication partner.

Configuring Replication Partners

WINS servers can use push, pull or both types of replication.

To add a replication partner to a WINS server, run the **WINS** MMC snap-in, right-click **Replication Partners** and choose **New Replication Partner**. Enter the **name** or **IP address** of this server's replication partner and click **OK**. Right-click on the **new partner's icon** and choose **Properties**. On the Advanced tab make the selections that are listed in the "Replication Types" section for Push and Pull Replication and click **OK**.

Configuring Global Options

Administrators can set options that apply to all domain WINS servers installed in the future (but not ones that have already been configured). To configure these options, right-click on the **Replication Partners** folder and choose **Properties**. The following tabs can be configured:

- **General** – Choose to replicate not only with partners but with other servers.
- **Push Replication** – Enable or disable the starting of push replication at: service startup, and when an address changes. Also has settings for the number of changes in version ID before replication takes place, and whether or not to use persistent connections.

- **Pull Replication** – Enable or disable the starting of pull replication at startup, and enable or disable persistent connections. Also contains entries for the replication start time and interval and the number of retries.
- **Advanced** – Has a setting to block certain servers from replicating and a setting, to be used only on small networks due to its multicast nature, which enables **automatic partner configuration**. For automatic configuration, every WINS server announces its presence with multicasts, and, if a server is found without a push/pull partner, this server gets added into the replication list of an existing server. Multicast settings include a time interval and time to live (TTL).

Plan NetBIOS name resolution by using the Lmhosts file

Besides using WINS to resolve NetBIOS names you can use an **Lmhosts** file. This file resides on every Microsoft client computer and needs to be updated on that local machine with any NetBIOS name changes. Unlike the **HOSTS** file, where you have to go to a particular machine and make changes, you can configure Windows 2000 and newer clients to *import* a shared Lmhosts file from a network share. This is the best way to plan for NetBIOS name resolution if you are going to use an Lmhosts file. This way you only have to update the file once and configure the clients to access that file. Below in *Figure 8* is a screenshot of how an XP client can be configured to import that shared Lmhosts file.

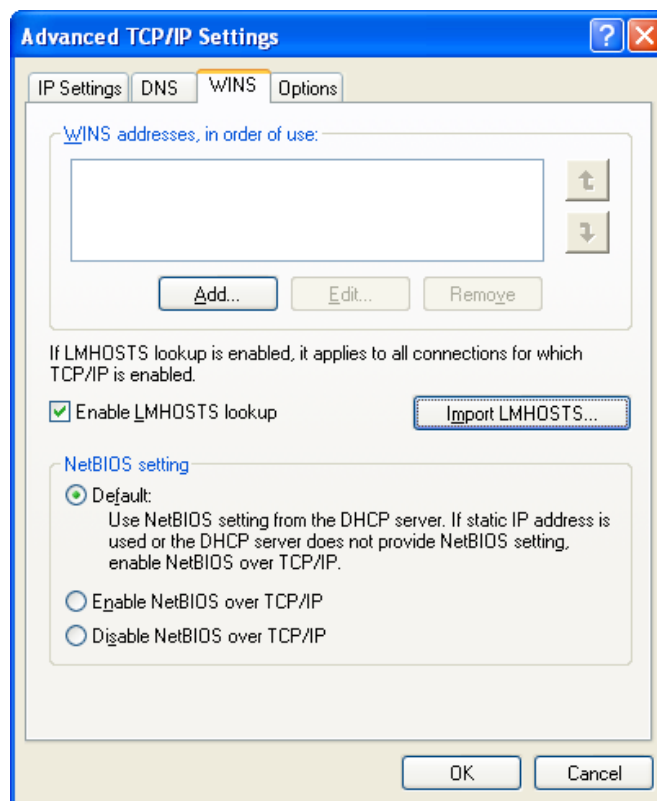


Figure 8 – Advanced TCP/IP Settings

Troubleshoot host name resolution

You need to know how to troubleshoot host name resolution problems for Windows 2000 clients and above using host name resolution as their primary name resolution method. First you will want to know the order of resolution used by a Windows 2000, XP and Server 2003 machine to resolve host names.

1. Is it me?
2. Is it in my local name cache?
3. Check HOSTS file.
4. Check DNS.

You can see that an incorrect entry in a client or server's HOSTS file can cause serious problems with your host name resolution. Therefore you should only use HOSTS files if needed; otherwise, stick with DNS for your host name resolution.

Diagnose and resolve issues related to DNS services

DNS servers are critical to a network; therefore, you must know how to manage, maintain and troubleshoot them to ensure that they are functioning properly and to optimize network performance. You can use the Event Viewer as a primary troubleshooting tool to view DNS events. This information is useful when you need detailed information on how the service is performing.

You can also set up debugging and enable detailed logging on your DNS server for advanced troubleshooting. To enable this go to the DNS Server's properties in the DNS console, and then click on the logging tab and enable which items you want to debug.

Depending on what type of zones you are using, a secondary DNS server may have issues as well. A secondary DNS server queries its primary DNS server for updates to a zone file and uses the serial number in the **SOA resource record** to determine whether changes have been made to the zone. If the serial number has changed, a zone transfer occurs to update the records on the secondary server. When a secondary server is not receiving updates from its master server, you can use the **Nslookup** utility to compare the serial numbers in each server's SOA resource record.

Diagnose and resolve issues related to client computer configuration

Most issues with host name resolution will reside on the client computers, and unfortunately it happens most often when a user attempts to configure something they read about on the Internet. The first step is to always check the Physical Layer. This simply means, "Are you plugged in?" More than half of your problems can be resolved with that simple question.

The next step to resolve configuration problems is to see what settings the client machine has configured for TCP/IP. This can be accomplished by using **ipconfig/all** command from a Command Prompt. This command will display every resource TCP/IP setting and option, and thus allow you to see if they are configured correctly for network communication.

Planning, Implementing, and Maintaining Routing and Remote Access

Plan a routing strategy

Routing is a method of transporting a data packet from a source to a destination. Of course, routing is not needed in cases where messages are destined for the same network subnet.

Windows Server 2003, by means of built-in utilities and the **Routing and Remote Access Service (RRAS)** – essentially a *multiprotocol router*, handles three types of routing:

- **Static** – Relies on a static routing table, which is stored in RAM and contains instructions on how to get a packet from one network segment to another. A routing table is created when TCP/IP starts up and changes when an administrator manually adds, or removes records.
- **Dynamic** – Relies on either **Distance Vector** or **Link State** routing protocols. The main benefit of dynamic routing over static is that records are added and changed automatically in the routing table (hence the name *dynamic*). This automatic operation leads to a related benefit: **fault tolerance**. If one network path goes down, the dynamic routing protocol will attempt to find an alternate route. One of the key ways to compare dynamic protocols is to see how they rate in terms of **convergence**: the elapsed time for all routers to record a routing information update.
- **Demand-Dial** – Also known as **Dial Demand Router (DDR)**, this method relies on a modem, ISDN or direct (serial or parallel) connection. This type of routing serves two main purposes: to backup a main connection for **redundancy**, and to **save on connection costs** (costs accrue only while the connection is in progress).

Identify routing protocols to use in a specified environment

Windows Server 2003 provides support for three routing protocols: NAT/basic firewall for small networks; RIP (**Routing Information Protocol**) for small to medium sized companies; and OSPF (**Open Shortest Path First**) for medium to larger sized companies.

NAT (Network Address Translation) is commonly used in small networks where RRAS is configured with two NICs – one for the LAN, the other for the Internet (WAN). The WAN NIC is used in place of a hardware Router. Similar in function to a hardware router, NAT secures a small network by blocking all incoming ports and, by default, only allowing basic Internet/e-mail connectivity.

To setup and configure NAT/basic firewall:

- Run **RRAS**
- Right-click **General** under the **IP Routing** folder
- Choose **New Routing Protocol**
- Select **NAT/Basic Firewall** Run **RRAS**
- Right-click **General** under the **IP Routing** folder
- Choose **New Routing Protocol**
- Select **RIP Version 2 for Internet Protocol** As shown in *Figure 9*

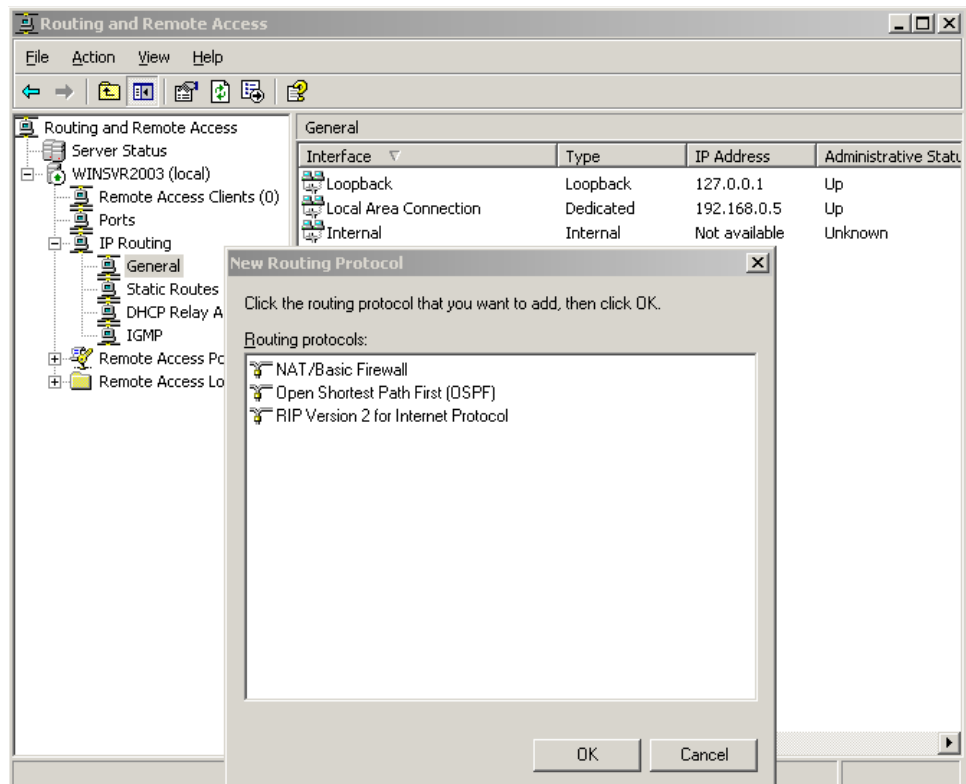


Figure 9 - NAT/Basic Firewall

RIP (Routing Information Protocol) routing is a **Distance Vector** style of routing that leverages the information it learns from other routers to build a routing table, and that uses a **hop count** metric (which normally increments by one as each new network segment is encountered) to count the number of routers which must be traversed to reach the destination network. RIP is restricted to a **maximum of 15 hops** and can send only **up to 25 routes in a single RIP packet**.

To add RIP routing, do this:

- Run **RRAS**
- Right-click **General** under the **IP Routing** folder
- Choose **New Routing Protocol**
- Select **RIP Version 2 for Internet Protocol**

Next, it's necessary to assign an interface to the RIP protocol. This is done in RRAS by right-clicking the **RIP** item and choosing **New Interface**. Select the **interface** that is attached to a network for which you want to learn more about routes.

In Distance Vector routing, routers update their record based on these events:

- When routers are started
- When routers have changes to their routing tables
- On a periodic basis

Advantages of Distance-Vector routing:

- It is **simpler to configure** than Link State
- It is **simpler to maintain** than Link State

Disadvantages of Distance-Vector routing:

- It is **slower to converge** than Link State
- It is at risk to the **Count-to-infinity** problem
- It **creates more traffic** than Link State since a hop count change must be propagated to all routers, and then the change must be processed on each router. Moreover, because of its periodic hop count updates, even if there are no changes in the network topology, broadcasts still occur thus wasting bandwidth
- For larger networks, it results in **larger routing tables** than Link State since each router must know about all other routers. This can also lead to congestion on WAN links

Note: RIP announces host or default routes by default.

The **Count-to-infinity** problem results when one router cannot reach an adjoining network. A second router, 1 hop away from the first router, thinks that the unreachable network is 2 hops away. Meanwhile, the first router then updates its records to say it is 3 hops away from the unreachable network based on the fact it is 1 hop from the second router, which says it is 2 hops from the unreachable network. The routers continue incrementing their hop count until the maximum (15), "infinity," is reached. There are three methods to prevent this problem: **Split Horizon**, **Split Horizon with Poison Reverse**, and **Triggered updates**.

RIP

RIP, which the Internet uses extensively, comes in two flavors: Versions 1 and Version 2.

RIP Version 1

To set up RIP Version 1, perform the same steps listed under the section "Distance Vector (RIP) Routing" and then add these steps before closing RRAS:

- Right-click the **RIP** icon and select **New Interface**
- Choose the **Local Area Connection** entry associated with the interface you wish to adjust and click **OK**

Operation Mode entries included **Auto-Static** update mode or **Periodic** update mode (the default). **Outgoing packet protocol** includes:

- RIP Version 1 Broadcast
- RIP Version 2 Broadcast
- Multicast
- Silent RIP

Incoming packet protocol includes:

- RIP Version 1 and 2
- RIP Version 1 only
- RIP Version 2 only

The **Security** tab has these settings for **incoming** and **outgoing** routes:

- Accept all routes
- Accept all routes in the ranges listed
- Ignore all routes in the ranges listed

The **Neighbors** tab allows the following settings for neighboring router addresses:

- Use broadcast or multicast only
- Use neighbors in addition to broadcast or multicast
- Use neighbors instead of broadcast or multicast; i.e., unicast

RIP Properties

RIP Properties include these settings (defaults are shown in brackets): Periodic announcement interval (30s); Time before route expires (180s); and Time before route is removed (120).

RIP Version 1 Shortcomings

One downfall of RIP Version 1 is that it **broadcasts** its route updates, and hence all hosts, not just RIP ones, get the announcements. However, this disadvantage of creating extra traffic has an advantage: a host can be configured as a **Silent RIP** and help with routing.

Silent RIP refers to a router that handles RIP announcements, while at the same time not announcing its own routes.

To configure Silent RIP, you would follow the same steps listed in this Exam Manual to configure RIP Version 1 but instead set the **Outgoing packet protocol** to **Silent RIP**.

RIP Version 1 *assumes a class-based networking scheme*. It thus assumes that Class A, B and C addresses use the default subnet masks (e.g., 255.0.0.0, 255.255.0.0 and 255.255.255.0). For addresses that don't map directly to Class A, B or C, if the network portion of the address matches that defined by the subnet mask associated with the interface sending the data, the system assumes that subnet mask, otherwise the 255.255.255.255 subnet mask is assumed. Because of these assumptions, supernetted routes may appear as a single network address, and a subnet being announced outside of its network segment may appear as a single host. As a result RIP Version 1 does not support **CIDR** and **VLSM**.

Another disadvantage of using RIP Version 1 is that it can't protect against a rogue RIP router from filling the network with garbage routes.

RIP Version 2

RIP Version 2 addresses some of the Version 1 shortcomings. It does this by:

- Using a **multicast** instead of a broadcast address. Instead of broadcasting RIP announcements, routers send to an IP multicast address of **224.0.0.9**. The disadvantage of this is that a Silent RIP router may not understand the multicast announcement.
- Sending not just an address but also a **subnet mask**. This overcomes RIP Version 1's problem of having to assume what the associated subnet mask is. This enhancement enables RIP Version 2 to support **VLSM** and **CIDR** and to allow for **non-contiguous network segments**, which RIP Version 1 cannot handle. Also, RIP Version 2 can increase the number of available subnets because it can support the all "0" and all "1"s subnets.
- Preventing rogue RIP routers, by supporting authentication methods such as passwords, and **Message Digest 5 (MD5)**.

OSPF

OSPF (**Open Short Path First**) routing is a **Link State** type of routing, which only sends out network topology information when changes have occurred. Routers learn about network topology, and hence build their routing tables, by sending out "**hello**" packets to adjoining routers. These packets also notify other routers that the sending router is **still alive**.

Also, **flooding** takes place, in which routers share route change information via **Link State Advertisements (LSAs)**. Since an LSA is not changed by any routers, but is forwarded to all routers, this information is considered **first-hand**. Also, an LSA contains information only about the originating router's neighbors and thus results in smaller routing tables.

An **LSA** consists of three elements:

1. A router
2. The networks attached to the router
3. The costs for the networks

OSPF also has these features:

- It is essentially **loop-free**, having a maximum hop metric of 65,535
- It can **load balance** network traffic between multiple paths of the same metric value
- It supports **authentication** using passwords and other methods

- It **converges quicker** than RIP since routing updates are sent immediately instead of periodically.
- It **uses less bandwidth** since transmission take place only when routing changes occur.
- It supports the **logical grouping of network segments** into areas (see the “Autonomous System” section below). Moreover, it announces routes outside of an autonomous system within the autonomous system so that it can calculate costs to reach outside networks.
- Thanks to announcing subnet masks, it supports **CIDR**, **VLSM** (Variable Length Subnetting), **Supernetting** (used to aggregate Class C networks) and **non-contiguous** network segments.

To add OSPF routing, follow the same steps as to add RIP except choose **Open Shortest Path First** instead of RIP.

The three main components of a routing hierarchy are:

1. **Autonomous System** – Networks with the same routing scheme and administration
2. **Areas** – Grouping of networks within an autonomous system. Areas are identified by an Area ID and are used to lessen the size of the **Topological Database**, which uses LSAs to provide a bird’s eye view of the networks and their proximity to routers. Each area forms its own topological database. Routers with more than one interface can link areas and are called **border routers**
3. **OSPF Backbone** – A special type of OSPF area, which links areas by transmitting routing information between them. Each OSPF internetwork has at least one backbone. The area ID for a backbone is **0.0.0.0** and is reserved. The backbone must be in the middle of all areas within an autonomous system and consists of:
 - ▶ Border routers
 - ▶ Networks not assigned to an area
 - ▶ Unassigned networks’ routers

A **Virtual Link** is used to connect an area, which cannot physically interface with an OSPF Backbone, to an OSPF Backbone. The Virtual Link is the connection between two routers, while the area that hosts those two routers is called a **Transit Area**.

Plan routing for IP multicast traffic

You use a multicast scope to issue a multicast address to selected computers on a network. When you use DHCP to configure client computers with a multicast address, these clients can participate in *collaborative application* sessions. Typically audio and conferencing applications, such as Microsoft Windows Media, use multicasting technology for deploying information from a single point to multiple computers at one time. You can configure several computers with the same multicast address in addition to each computer’s individual IP address. All computers configured with the same multicast address receive IP packets that are sent to that address. For multicasting to work correctly, all routers between the server that is sending packets to the multicast address and the receiving client computers must be configured to recognize the multicast address.

Configuring a multicast scope to issue a multicast address eliminates the need for users to have to specify the address manually. To take advantage of dynamic multicast IP addressing, you must configure a multicast scope on a DHCP server, and multicast applications on client computers must be able to use the **Multicast Address Dynamic Client Allocation Protocol (MADCAP)**.

The range of addresses that can be used for multicasting are 224.0.0.0 to 239.255.255.255.

Plan security for remote access users

In Windows Server 2003, the **Routing and Remote Access Service (RRAS)** is installed automatically, though not activated. This service includes the RAS (Remote Access Service) functionality available under Windows NT 4.0 as well as VPN (**Virtual Private Networking**) and other enhancements. Windows Server 2003 uses **Remote Access Policies** to help secure the incoming connections to a network.

Plan remote access policies

Remote access policies are a set of rules that determine whether a user will get a connection or be rejected. You can use multiple remote access policies to control the access of different users and user groups. If you use multiple remote access policies, then policies will be checked from top to bottom. You can reorder the policies for a desired result. You should order the policies beginning with the most specific policies at the top of the list. Once a policy matches, it will be used and no further policies will be checked.

Remote access policies are much more than just permissions to dial in to a network. In fact, they contain three components that work together to accept or deny a connection to the network. You can configure each of these components to achieve your desired result. You should know how to configure these components of remote access policy:

- Conditions
- Permissions
- Profile

Remote access policy **conditions** are attributes that must be met in order to satisfy the policy. Conditions are only checked at the initial time of the connection attempt and are thus the first component that is checked on a connection attempt. Conditions might include day and time restrictions, connection types, security group memberships, and many others. If you set multiple conditions on the same remote access policy, then all of the conditions have to be met.

The Dial-in **permissions** of the user are checked after the conditions, assuming that a condition to deny has not already been met. If your domain is in at least the Windows 2000 native mode functional level, then the Dial-in permissions for the user can be set to **Allow**, **Deny**, or **Control Access through Remote Access Policy**. If your domain is in a lower functional level, then the Control access through Remote Access Policy option is not available. If you set the permissions to Allow, then the user is connected because they already met the conditions earlier. If you set the permissions to Deny, then the user is denied access, even though they met the conditions earlier. If you set the Permissions to Control access through Remote Access Policy, then the user's connection will be accepted or denied based on the next step.

If you set the user permissions to Control access through Remote Access Policy, then the **profile** settings on the policy must be met in order to obtain and to continue a connection. Profile settings that you can select include day and time restrictions, idle-timeouts, session-timeouts, encryption, authentication, connection types, and many more. If you set multiple profile settings in a remote access policy, then the user must meet and continue to meet the restrictions that you set.

Analyze protocol security requirements

When working with RRAS, you might find it helpful to keep in mind that all connections, including ones made via modem, are treated like LAN connections. The advantage of this is that full Windows 2000 network functionality is available to dial-in clients; the disadvantage is that this same functionality must be carefully secured.

RRAS supports these networking protocols: **AppleTalk, IPX, NetBEUI, TCP/IP.**

It also supports these data link control protocols for asynchronous connections:

- **SLIP (Serial Line Interface Protocol)** – Older, limited protocol for use only with legacy applications. It can be used only by clients
- **PPP (Point-to-Point Protocol)** – Enhancements over SLIP include: error correction, compatibility with several authentication protocols, support for more protocols than just TCP/IP, and automatic session setup and disconnection

Plan authentication methods for remote access clients

Remote access servers use authentication to determine the identity of users attempting to connect to the network remotely. After a user is authenticated, the user receives the appropriate access permissions and is allowed to connect to the network. The correct and secure authentication of user accounts is critical for the security of a network. Without authentication, a potentially large number of unauthorized users could access your network.

Routing and Remote Access uses several protocols to perform authentication, and also allows for the use of **Extensible Authentication Protocols (EAP)**, through which you can load third-party protocols.

Protocol	Security	When to Use
PAP	Low	Client and server cannot negotiate using more secure validation
SPAP	Medium	Using Shiva LANRover or a Shiva client
CHAP	Medium	Clients that are not running a Microsoft OS
MS-CHAP	High	Clients running Windows NT, 95 and later
MS-CHAP v2	High	Dial-up clients running Windows 2000 and above, VPN clients running Windows NT or 98 and above

EAP allows for customized authentication to remote access servers. The client and the remote access server negotiate the exact authentication method to be used.

Remote Authentication Dial-In User Service (RADIUS) is an industry-standard protocol that provides the solution to these authentication and remote user accounting requirements. The combination of Routing and Remote Access and the **Internet Authentication Service (IAS)** provides support for RADIUS. The following steps describe the basic process that remote servers, a **RADIUS** server, and RADIUS clients use to perform authentication and authorization:

1. A user connects to a Server 2003 based computer that is running Routing and Remote Access by using a dial-up connection or a VPN connection.
2. The Windows Server 2003 based computer that is running Routing and Remote Access forwards authentication requests to an IAS server. When doing this, the computer running Routing and Remote Access acts as a RADIUS client.
3. The IAS server accesses the user account information on a domain controller and checks the remote access authentication credentials. When doing this, the IAS server performs the functions of a RADIUS server.
4. If the user's credentials are authenticated and the connection attempt is authorized, the IAS server authorizes the user's access and logs the remote access connections as accounting events.

Implement secure access between private networks

It is very rare that you will see a network that does not communicate with other networks. Because this is such common practice you must ensure that the network traffic going between networks is secure. Microsoft has adapted an industry standard protocol to help with the secure communication of IP traffic. The name of the protocol is **IPSec** (IP Security) and this has been included with Windows since Windows 2000. Microsoft has taken an already rock solid protocol and fine tuned it even more in Windows Server 2003 to optimize it for better performance and security.

IPSec uses **Kerberos V5** as its default method for authentication, but also supports **pre-shared keys** and **public key certificates** (X.509). No data exchange is allowed until both communicating computers authenticate with each other. IPSec performs the encryption for **L2TP** (Layer Two Tunneling Protocol).

Create and implement an IPSec policy

IPSec policies can be created and applied on local machines or through Group Policy in an Active Directory Domain. There can only be one active policy per machine at any given time. Microsoft includes three built in policies that should be used as a guideline, but never to be used in an actual network environment until they are modified to suit that environment.

- **Client (Respond Only)** – Security is used only if requested
- **Server (Request Security)** – All sessions request IPSec security, but communications will proceed without it
- **Secure Server (Require Security)** – IPSec Security must be used

The really interesting thing about the above policies is that any of them can be applied to any type of machine. For example, a Windows Server 2003 machine could have the Client policy applied, while a Windows XP Professional machine could have the Secure Server policy. You will want to remember what the policies do and not so much the names of them.

Troubleshoot TCP/IP routing; Tools might include the route, tracert, ping, pathping, and netsh commands and Network Monitor

Windows Server 2003 includes many tools that will help you troubleshoot potential routing problems. These tools may be tools that you have had experience with in Windows 2000.

Command	Brief Description
Route	Displays and modifies the local routing table
Tracert	Determines the path to a destination by pinging each stop and incrementing the TTL (Time To Live)
Ping	Verifies IP connectivity between hosts
Pathping	A combination of tracert and ping, this determines the path to a destination and then, by default, pings each stop 100 times
Netsh	Utility that allows you to display or modify network configurations locally or remotely

Use **Network Monitor** to monitor the network data stream (packets), which consists of all information transferred over a network at any given time. The components of Network Monitor that you must install are:

- Network Monitor drivers, which capture the data
- Network Monitor tools, which allow you to view the data

Planning, Implementing, and Maintaining Server Availability

Plan services for high availability

Most companies need to have certain servers and services that maintain high availability. This simply means that the server and/or service is relying on either more reliable hardware and/or built in services to keep it from becoming a single point of failure. Windows Server 2003 contains support for just that. However you will find that for the most part these features are only available on high-end operating systems, such as **Windows Server 2003 Enterprise and Windows Server 2003 Datacenter server**.

Plan a high availability solution that uses clustering services

The cluster configuration process has been completely rewritten for Windows Server 2003 to eliminate previous shortcomings and to provide more detailed logging during the configuration, which allow users to fix problems before systems are fully in production and require expensive downtime to fix.

Clustering services can only be installed on a Windows Server 2003 Enterprise or Datacenter server. Clustering Windows Servers 2003 provides fault tolerance by using Failover and Fallback options to maintain continuous uptime for critical applications. The most overlooked issue with clustering is not having compatible hardware. When planning for clustering, always check the **HCL (Hardware Compatibility List)**. Another pitfall that many administrators fall into is that clustering does not support Dynamic Disks. Microsoft recommends that you have two network adapters although one can be used for cluster servers.

Plan a high availability solution that uses Network Load Balancing

Network load balancing (NLB) has been around for some time but Microsoft has added some very useful features to help when planning and administrating a Windows Server 2003 load-balancing configuration. A new GUI-based management tool, the Network Load Balancing Manger, allows NLB cluster management. The tool is not MMC-based, but has a similar layout. It provides a GUI for cluster creation, host addition and removal, cluster management, host management, and cluster deletion. It provides for rudimentary error checking, unlike the previous CLI-based management tool. It also provides for management of some of the new NLB features, including:

- Multi-NIC
 - NLB now can be bound to multiple NICs on a node. Each NIC must belong to a different cluster. Each of those clusters can co-exist on a host and can be managed independently.
- Virtual Cluster
 - Virtual clusters are used to manage different independent port rules. **Multi-VIP** addresses are assigned to an NLB cluster. Each VIP has an associated port rule that can be managed separately. Virtual cluster simplifies management of large and complex NLB clusters by separating control of the port rules. Each VIP can be managed independently.
- IGMP Support
 - NLB can induce “switch flooding” when multiple cluster hosts are connected directly down-stream from a switch. Switch flooding occurs because the switch cannot establish a 1-to-1 relation between the cluster MAC and the switch port, so it forwards the NLB cluster to all ports. In Windows Server 2003, NLB can use the **IGMP** protocol to in-

form the switch which ports belong to the cluster, thus preventing forwarding cluster traffic to ports that are not associated with the cluster.

Identify system bottlenecks, including memory, processor, disk, and network related bottlenecks.

The hardware in your server is of critical importance when purchasing a new server. If you decide to save a few bucks here and there for server hardware, then you may run into a system bottleneck. You will always want to keep a baseline of your system resources so that you can compare your currently monitored data to your historical baseline. This will assist you in troubleshooting for system bottlenecks. Microsoft has identified four key resources to monitor for system bottlenecks. System Monitor, described earlier in this Exam Manual, is the preferred tool to use for troubleshooting system bottlenecks. In the table below, the memory, CPU, and hard disk resources show the major counters and their descriptions.

Resource	Counters	Description
Memory	Pages/Sec	Indicates the number of requested pages that were not immediately available in RAM and had to be read from the disk or written to the disk to make room in RAM for other pages.
	Available Bytes	Indicates how much physical memory is remaining after the working sets of running processes and the cache have been served.
CPU	% Processor Time	Shows the percentage of elapsed time that a processor is busy executing a non-idle thread. Constant value of 85% or more indicates a need for a faster or additional processor.
	Processor Queue Length	Shows the count of threads currently in the processor queue. A queue of two or more items on a single-processor system may indicate a bottleneck.
Disks	% Free Space	Reports the percentage of unallocated disk space to the total usable space on the logical volume.
	Avg. Disk Bytes/Transfer	Measures the size of I/O operations. The disk is efficient if it transfers large amounts of data relatively quickly.
	Avg. Disk sec/Transfer	Indicates how fast data is being moved (in seconds). Measures the average time of each data transfer, regardless of the number of bytes read or written.
	Disk Bytes/sec	Indicates the rate at which bytes are transferred and is the primary measure of disk throughput. The higher the number the better the performance.
	Disk Transfers/sec	Indicates the number of read and writes completed per second, regardless of how much data they involve. Measures disk utilization.
Network	Output Queue Length	Indicates the length of the output packet queue. The value should be low.

	Packets Out-bound Discarded	Use this counter to determine if the network is saturated.
	Bytes Total/sec	Use this counter to monitor the performance of the network adapter. High values indicate a large number of successful transmissions.

Identify system bottlenecks by using System Monitor

You can use System Monitor to obtain more comprehensive information about your computer or other computers on the network. This is the premier tool that comes with Windows Server 2003 that allows you to identify system bottlenecks. You can use information collected by System Monitor to diagnose how the system and applications are functioning, and to ensure that you are optimizing your server system. For example, if your network is running slower than normal, you can use System Monitor to monitor the CPU percent processor time.

System Monitor information can be viewed in graph, histogram (bar chart), or report form. Graphs, histograms, and reports can also be viewed in a browser and printed when performance data is saved as a Hypertext Markup Language (HTML) file.

Implement a cluster server

When creating a cluster, you should always start with the New Server Cluster Wizard, as shown in *Figure 10*.



Figure 10 – Implementing a Cluster Wizard

This wizard walks you through setting up a cluster server and also helps with cluster planning. The steps for setting up a cluster server are depicted in *Figure 10*.

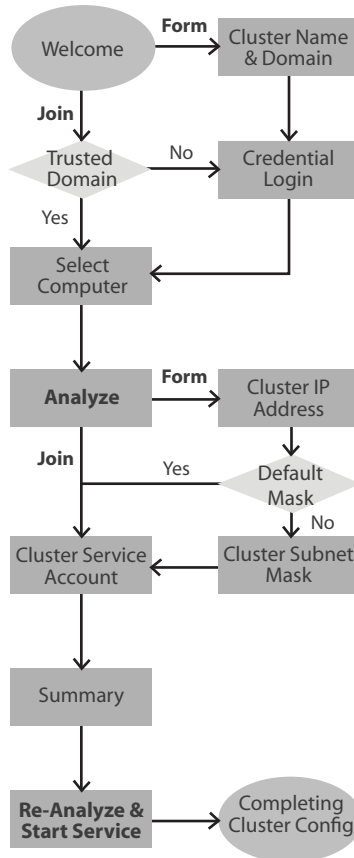


Figure 11 – Steps for Setting up a Cluster Server

Recover from cluster node failure

There are several ways to recover a failed cluster in Windows Server 2003. You need to understand what tools you should use to recover in each instance.

- Loss of a Shared Cluster Disk
 - Loss of a shared disk in Windows Server 2003 clusters is much simpler and more automated than in previous Windows operating systems. There are two preferred methods for recovering from a lost shared disk: [Confdisk.exe](#) and [Clusterrecovery.exe](#), both residing in the Windows Server 2003 Resource Kit CD-ROM. Of these two tools, Confdisk is the preferred method because it is more automatic.

- Loss of a Cluster Node
 - The loss of a cluster node in Windows Server 2003 can be recovered using **ASR (Automated System Recovery)** in combination with NTBackup **System State** restore. ASR will restore the operating system to the node as well as the disk signatures and partition layouts of critical disks, such as the local system and boot partitions. Applications and local data partitions that are not recovered in the process of ASR should be recovered using a restore from a full backup once the ASR process is completed and the operating system is restarted.
- Loss of the Entire Cluster
 - If the entire Windows Server 2003 cluster is lost or upgraded to new hardware, restoration would involve rebuilding one of the nodes with **ASR**, and then performing a **System State** restore and an **Authoritative** restore of the cluster state. Since only one node is running at this point, it will have access to the shared disks, so ASR will rebuild the disk signature and partition information on these disks. After the first node is completed the remaining nodes would be rebuilt using ASR and would rejoin successfully after the ASR rebuild is completed and the systems are rebooted.
- Unresolvable Configuration Problems
 - In Windows Server 2003 cluster, an undesirable cluster state can be brought back to a known good state using a previous system state backup and choosing **Restore the cluster registry to the quorum disk and all other nodes** in the **Advanced** settings of NTBackup. This feature will restore the cluster state to all nodes in the cluster simultaneously. Note that this procedure does take down the entire cluster and it restores all of the system state information to the node that you run it on.

Manage Network Load Balancing. Tools might include the Network Load Balancing Monitor, Microsoft Management Console (MMC) snap-in, and the WLBS cluster control utility

Two tools that help you manage your clusters and Load Balancing services are the **NLB Monitor** and **WLBS (Windows Load Balancing Service) Cluster Control Utility**. The look and feel of NLB Monitor is a tool that will be very familiar to administrators as it snaps-in into an MMC. The WLBS Cluster control utility provides the ability to manage not only Windows Server 2003 clusters, but also other clusters running on older Microsoft clients from a command prompt.

Plan a backup and recovery strategy

Backing up data is one of the most important job functions for any administrator. Administrators often make a common oversight when it comes to a backup and recovery strategy: they don't test the recovery strategy. Test restores should be periodically performed to assure future data recovery. You should be able to restore your data with minimal downtime, as well as ensure the backed up media data integrity.

Microsoft has improved the backup software that comes with Windows Server 2003. They combined popular features from Windows 2000 and XP and even added an additional feature that will save tech support a lot of time when restoring deleted files.

Identify appropriate backup types. Methods include full, incremental, and differential

Choosing the most appropriate backup type is an important factor that many administrators often overlook. Companies will have different needs and administrators should consider each company's backup strategy on a case-by-case basis. Remember, Differential backups take the longest time to backup and the shortest time to restore, while Incremental backups backup data faster than Differential ones but take much longer to restore. Below you see the different backup methods and additional info about each.

Type	Backs up	Clears marker
Normal	Selected files and folders	Yes
Copy	Selected files and folders	No
Differential	Selected files and folders that changed since the last backup	No
Incremental	Selected files and folders that changed since the last backup	Yes
Daily	Selected files and folders that changed during the day	No

Plan a backup strategy that uses volume shadow copy

Shadow Copies for Shared Files uses the **Volume Shadow Copy Service (VSS)** found in Windows Server 2003 to create point-in-time changes to files and folders located on network shared folder volumes for which VSS is enabled. Because VSS is used, Shadow Copies occur without service interruption. Shadow Copies for Shared Files provides an end user-accessible tool that can restore, copy, or view previous versions of users' documents by accessing point-in-time shadow copies of documents and folders stored on the shadow copy volume. Shadow Copies for Shared Files is transparent to end users when they create, save, open or delete files on network file shares. Users no longer have to endure lost productivity while they wait for the IT department to restore a single file from tape backup. IT organizations who implement Shadow Copies for Shared Folders for file recovery scenarios realize a number of benefits, including: reduced demand on busy administrators through the reduction in restore-from-tape requests; and reduced cost of recovering single or multiple files or folders.

The default schedule for Shadow Copies is to take point-in-time snapshots twice daily, Monday through Friday at 7AM and Noon. The default size of the Shadow Copy Diff area is 10% of the volume (100 MB minimum). You can adjust both the Time Schedule and Shadow Copy Diff area by clicking the Settings button on the Servers disk volume Properties, Shadow Copy tab as shown in *Figure 12*. When the Diff area fills, the oldest Shadow Copy is dropped off to make room for the current Shadow Copy. For this reason, Shadow Copies of Shared Folders should not be used in place of archival systems.

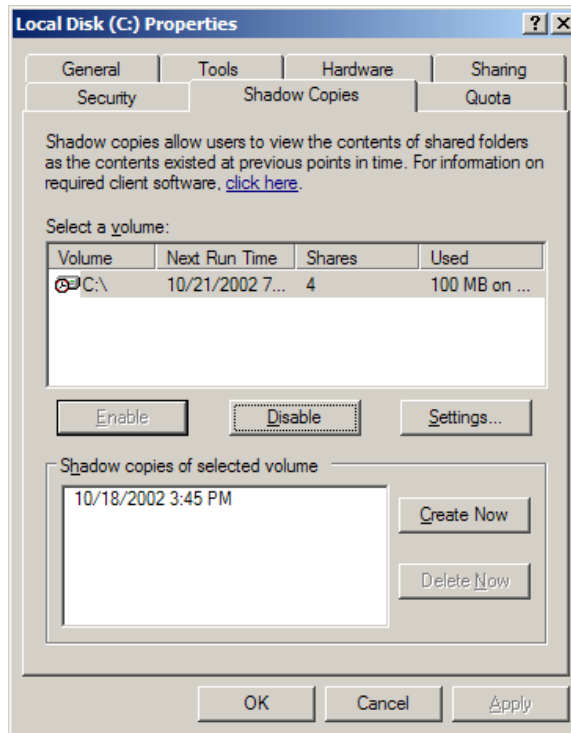


Figure 12 – Configuring Shadow Copies

Only Windows Server 2003 and Windows XP Professional SP1 come with the Shadow Copy client feature installed. To install Shadow Copy Client for earlier Windows versions of Windows XP Professional and Windows 2000 (service pack 3 or later) operating systems, you first have to download the Shadow Client install file from Microsoft's download site. Once installed, the Previous Versions tab displays previous versions of shared folders and files as shown in *Figure 13*.

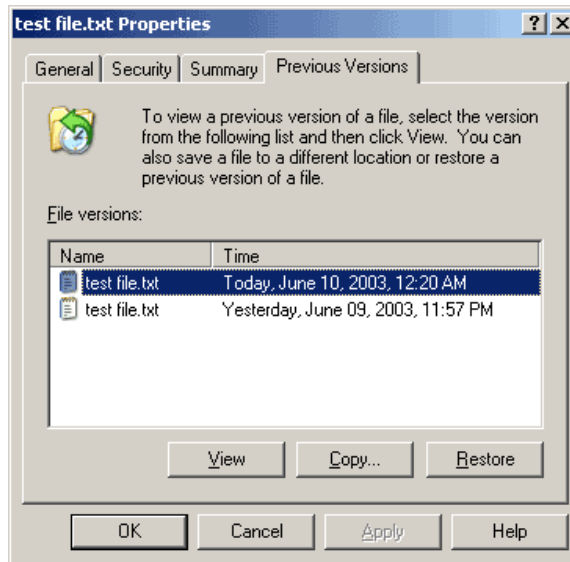


Figure 13 – Viewing Previous Versions of a File

Plan system recovery that uses Automated System Recovery (ASR)

The Automated System Recovery (**ASR**) feature provides the ability to recover Windows Server 2003 operating system from a Server hardware failure or corrupted operating system. ASR creates an image file of your operating system and stores necessary startup/recovery information on a floppy disk that you insert when prompted during the ASR backup process. ASR has two parts, **ASR Backup** and **ASR Restore**. It is used to back up the server operating system; however, it does not include data files. ASR is not intended to back up applications. When you have tried all other recovery methods like Last Known Good Config, Driver Rollback, Safe Mode, etc. then use ASR to recover your Windows Server 2003 OS, system state and Registry. Use ASR as a last resort. To backup ASR, start the backup utility, select Advanced mode and then select the Automated system Recovery Wizard from the Welcome window as shown in *Figure 14*. To restore an ASR backup, access ASR by booting from your Windows Server 2003 CD-ROM and pressing F2 at startup to invoke the ASR Wizard.

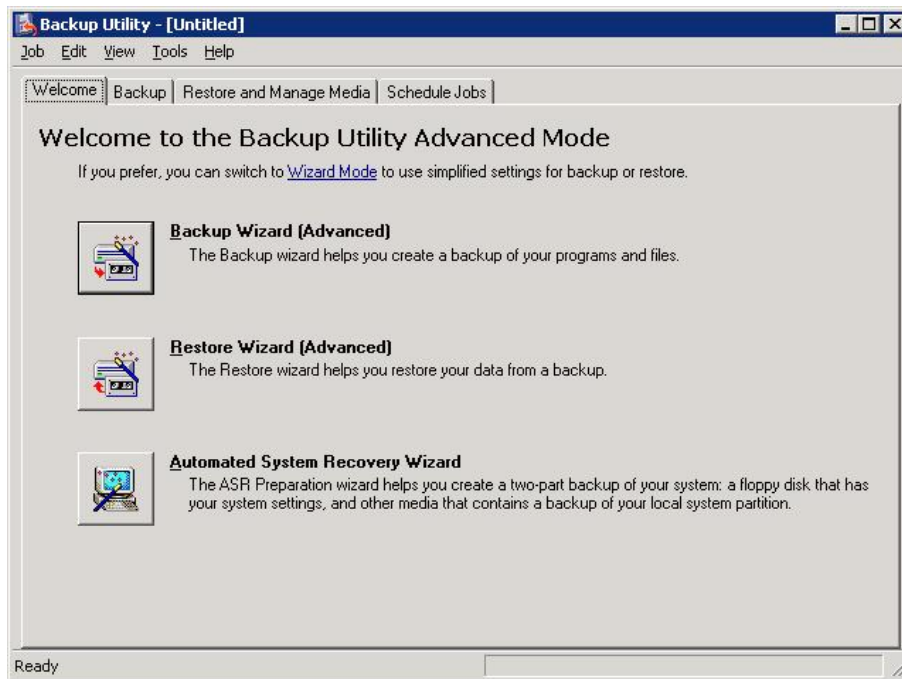


Figure 14 – Accessing the Backup and Restore ASR Wizards

Planning and Maintaining Network Security

Configure network protocol security

Before you can configure network protocol security you must determine which protocols are being used on your network. Once you have determined which protocols are in use you can then begin configuring security. Don't expect because all your clients are running the same operating system that you can blanket them all with the appropriate security. There are several factors that will decide what type of security will be configured. Some of these factors stem from the application and business use of these machines.

One thing is certain: if you have to configure security on any Windows 2000 and above machine, you will want to use Group Policy. Group Policy will give you the granular control that you need when multiple computers in different departments have different protocol security needs.

Configure protocol security in a heterogeneous client computer environment

When working in a heterogeneous network environment it is critical to know which protocols client computers will use. You may have Novell, UNIX or even Apple machines in your network at any given time. If you have older Novell clients or servers on your network, then you may need to use IPX/SPX, or NWLink. NWLink has no built-in security features and should only be used if necessary. Recent Novell's NetWare servers use TCP/IP and contain support for IPSec as well.

If you have UNIX machines in your network you will need to be aware of several services that are not secure. FTP is a great file transfer protocol but the protocol data is sent across the network in clear text. This also goes for the authentication traffic. To secure FTP traffic you should use IPSec since FTP uses TCP/IP. **SNMP (Simple Network Management Protocol)** is another protocol used for troubleshooting networks that is utilized on many types of network. If you use SNMP, make sure to always change the default community password from public to something more secure.

Finally, another service that communicates in clear text is **Telnet**. This can be used to configure UNIX and Linux machines and other networking devices. Thankfully, it runs on IP and IPSec can also be used to secure it.

Configure protocol security by using IPSec policies

IPSec is Microsoft's recommended security solution for securing networks. It provides strong, secure protection on your private network from Internet attacks through end-to-end security. The only computers that need to know about IPSec are the sending and receiving computers. For the Windows 2000 Professional and XP clients and the Windows Server 2003 family, IPSec provides the ability to protect communication between workgroups, local area network computers, domain clients and servers, branch offices, extranets, and roving clients.

IPSec is implemented primarily to enforce security policies for IP network traffic. A security policy is a set of rules that define network traffic at the IP layer. A packet filter action defines the security requirements for the network traffic. A filter action can be configured to:

- Permit
- Block
- Negotiate security (negotiate IPSec).

IPSec filters are inserted into the IP layer of the computer's TCP/IP networking protocol stack so that they can examine and filter all inbound or outbound IP packets. Except for a brief delay required to negotiate a security relationship between two computers, IPSec is transparent to end-user applications and operating system services.

To ensure that IPSec communication is successful and that IPSec meets the security requirements of your organization, you must carefully design, configure, coordinate, and manage IPSec policies.

IPSec policies are accessed and configured by using the Group Policy Object Editor and navigating to Computer Settings\Windows Settings\Security Settings\IPSec Policies on Active Directory. *Figure 15* displays the three types of IPSec policies in the right pane:

- **Server - Request Security:** Clients will request security using Kerberos trust. Allows unsecured communication for clients that are not configured or do not support Kerberos.
- **Client - Respond Only:** By default, clients will communicate normally unsecured. Clients will, however, use secured IPSec with servers that request security.

Secure - Requires Security: The most secure method that always requires security using a Kerberos trust. Does not allow unsecured communication and thus should not be used on networks containing down-level clients like Windows NT workstation or Windows 9.x

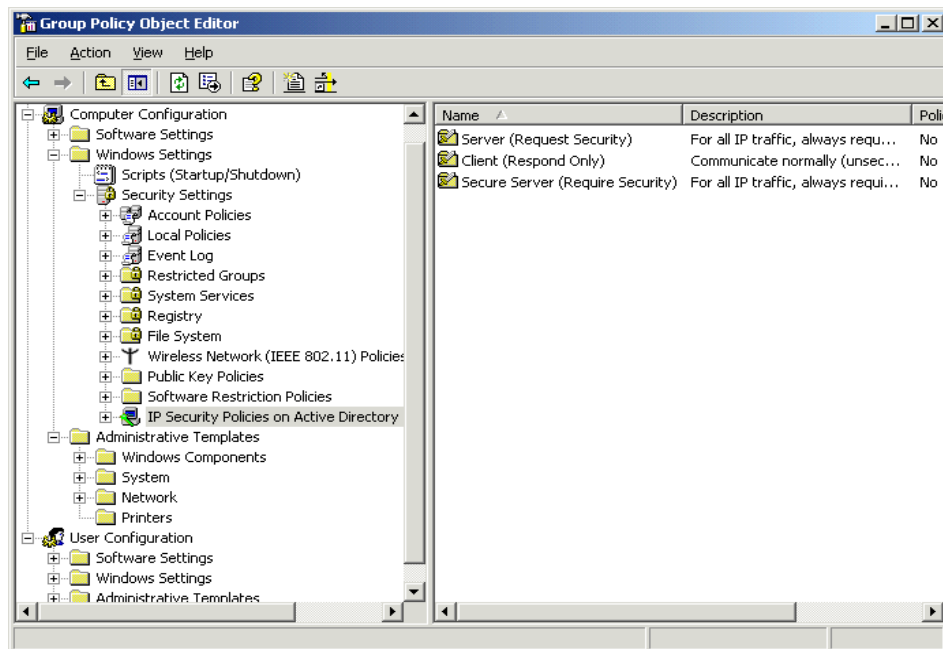


Figure 15 - Group Policy Object Editor highlighting IP Security Policies

Configure security for data transmission

In Windows 2000 and Windows XP, by default, broadcast, multicast, **Kerberos**, **Internet Key Exchange (IKE)**, and **Resource Reservation Protocol (RSVP)** traffic is exempt from filter matches. In the Windows Server 2003 family, broadcast, multicast, Kerberos, and RSVP traffic is not exempt from filter matches by default (only **IKE** traffic is exempt). Broadcast and multicast packets will be dropped if they match a filter with a filter action to negotiate security.

By default, the Windows Server 2003 family provides limited support for filtering broadcast and multicast traffic. A filter with a source address of Any IP Address will match multicast and broadcast addresses. A filter with a source address of any IP Address and a destination address of any IP Address will match inbound and outbound multicast addresses. You can use such a filter to block all traffic. One-way filters that would be used to block or permit specific multicast or broadcast traffic, however, are not supported.

Configure IPSec Policy Setting

As a result of the change in default exemption behavior for the Windows Server 2003 family implementation of IPSec, you should test and verify the behavior of IPSec policies designed for Windows 2000 or Windows XP clients, and determine whether to configure explicit permit filters to permit specific traffic types. To restore the default Windows 2000 and Windows XP behavior for IPSec policies, you can use the **netsh ipsec dynamic set config** command, or you can manually modify the appropriate registry settings.

Plan for network protocol security

Part of your initial design for your network infrastructure should include what types of protocols are going to be used. With that list you will want to do your homework and see if there is any security vulnerabilities associated with them. Remember that a lot of IP based services like **FTP**, **TELNET**, and **HTTP** all use clear text, by default. You will have to find an alternate way to secure that traffic if needed. The best way to do so would be to use an IPSec Security policy.

Specify the required ports and protocols for specified services

Servers run many different services. It is critical for you to know what ports and protocols these services run on. Once you identify the services and protocols on your network, you can then configure ports using IP filtering on your firewalls to allow specific types of traffic. A good rule of thumb is to initially deny everything and then open up only the ports you need. The table below lists popular server services and their associated ports and protocols.

Service	Port	Protocol
FTP-Data	20	TCP
FTP	21	TCP
Telnet	23	TCP
SMTP	25	TCP
DNS	53	TCP/UDP
HTTP	80	TCP
POP3	110	TCP
RPC Location	135	TCP
NetBIOS	137	TCP
NetBIOS Datagram	138	UDP
NetBIOS Session	139	TCP
IMAP	143	TCP
SNMP	161	TCP/UDP
SNMPTRAP	162	TCP/UDP
LDAP	389	TCP/UDP
HTTPS	443	TCP
SQL	1433	TCP

PPTP	1723	TCP
IAS (RADIUS)	1812	UDP
IAS (RADIUS)	1813	UDP
Remote Desktop	3389	TCP/UDP

Plan an IPSec policy for secure network communications

IPSec policies can secure network communication in two ways. It has the ability to **verify** the integrity of the sender and/or **encrypt** the entire communication between two hosts.

AH (Authentication Header) is an optional IPSEC protocol used to ensure *authenticity*. It uses the source and destination IP address as part of its integrity check. If the IP address or any data in the IP payload changes, then the integrity check fails.

ESP *encrypts* the IP payload and is typically the protocol used with **IPSEC tunnel** mode and the L2TP/IP-SEC (**transport** mode).

Plan secure network administration methods

Microsoft has included several network administrative methods with Windows Server 2003. You can remotely connect to Windows XP client machines through a technology called **Remote Assistance**. This feature should save your help desk support personnel time since they don't have to visit each client machine that has a problem. By default, Windows XP clients have the **Allow users to connect to this computer** check box enabled in the computer's system Properties Remote tab, Remote Desktop section. You can configure Remote Assistance manually by checking **the Allow Remote Assistance invitations to be sent to this computer** or via Group Policy. Windows Server 2003 Server also expanded on Windows 2000's **Terminal Services** by adding new features and a much richer client known as Remote Desktop Connection.

One service that is often overlooked is **Telnet**. As you will see, Remote Assistance and Remote Desktop work great for Microsoft clients and servers that support it, but what about the other devices and servers you may have, like routers, switches, UNIX servers and others? For other servers and network devices on your network you can use Telnet. The Telnet application is part of the TCP/IP protocol suite, and any network using TCP/IP can use it. Telnet is built into Windows Server 2003 and provides a command-line interface to remotely connect to another server like Unix or Linux for limited functionality, in order to configure or troubleshoot the server. Realize, though, that Telnet does not provide security and all passwords and data are transmitted in clear text. The table lists various Telnet commands and their actions performed.

Telnet Command	Action Performed
Open hostname	Establishes session with host
Close	Closes connection
Display	Shows current settings for client
Send	Gives additional commands as defined by type of host
Set	Used with additional arguments for configuration options; these are dependent on client
Unset	Turns off options that were previously set
Status	Determines connection status
?	Shows help menu based on host
Quit	Closes Telnet Client

Create a plan to offer Remote Assistance to client computers

Clients can request your assistance using the Remote Assistance tools included in Windows XP Professional and you can respond to their requests and assist them through your Windows Server 2003 **Remote Desktop Connection (RDC)**. After you are connected, you will be able to view the client's computer. You can even take control of their mouse and keyboard with their permission. You can also upload files to clients or download their files to your computer or a central server. Remote Assistance communication can also be used Windows Messenger or Microsoft Outlook.

Plan for remote administration by using Terminal Services

Using Windows Server 2003, **Remote Desktop Connection** replaces the Remote Administration Mode for Terminal Services used in Windows 2000. Windows Server 2003 provides a new interface that allows you to safely manage any computer that is configured to allow users to connect remotely. You can access Remote Desktop Connection through the Communications folder within Accessories on the Start menu. You then connect to the computer by entering the computer name and the password for that computer. You must also be a member of the **Remote Desktop Users** security group to use Remote Desktop Connection. The administrator is a member of this group by default.

Using RDC you can control the resolution and other aspects of the "user experience" on the Remote Desktop Connection settings. Other options allow you to configure your remote session based on the allowed bandwidth and other restrictions or use printers and drive mappings. You should use Remote Desktop Connection when you are making a connection to only one other computer or server. If you need to make multiple connections you should use the Remote Desktop MMC snap in with a custom MMC. *Figure 16* shows the Remote Desktop Connection Window displaying the Display tab properties.

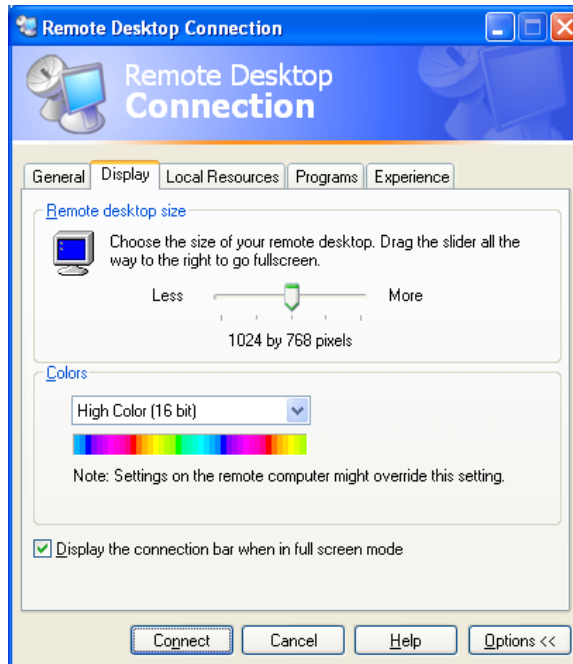


Figure 16 – Remote Desktop Connection Settings

Plan security for wireless networks

For authentication, the **IEEE 802.11** wireless standard defines open system and shared key authentication types. For data confidentiality, the 802.11 standard defines **Wired Equivalent Privacy (WEP)**. The 802.11 standard does not define or provide a WEP key management protocol that provides automatic encryption key determination and renewal.

This is a security limitation to IEEE 802.11 security services - especially for **wireless infrastructure mode** with a large number of wireless clients. This security issue is solved by using IEEE 802.1X port-based network access control.

IEEE 802.11 Authentication

Extensible Authentication Protocol-Transport Level Security (EAPTLS) is one type of authentication used for IEEE 802.11 networks. IEEE 802.11 defines the following types of authentication:

- Open System Authentication
- Shared Key Authentication

Open system authentication does not provide authentication, only identification using the wireless adapter's MAC address. Open system authentication is used when no authentication is required. Open system authentication is the default authentication algorithm that uses the following process:

1. The authentication-initiating wireless client sends an IEEE 802.11 authentication management frame that contains its identity.
2. The receiving wireless node checks the initiating station's identity and sends back an authentication verification frame.

With some wireless APs (**Access Points**), you can configure the MAC addresses of allowed wireless clients. However, this is not secure because the MAC address of a wireless client can be spoofed. By default, a Windows XP wireless client configured to perform open system authentication sends its MAC address as the identity.

Shared key authentication verifies that an authentication-initiating station has knowledge of a shared secret. This is similar to preshared key authentication for Internet Protocol security (IPSec). The 802.11 standard currently assumes that the shared secret is delivered to the participating wireless clients by means of a secure channel that is independent of IEEE 802.11. In practice, this secret is manually typed at the wireless AP and the wireless client.

Shared key authentication uses the following process:

1. The authentication-initiating wireless client sends a frame consisting of an identity assertion and a request for authentication.
2. The authenticating wireless node responds to the authentication-initiating wireless node with challenge text.
3. The authentication-initiating wireless node replies to the authenticating wireless node with the challenge text that is encrypted using **WEP** and an encryption key that is derived from the shared key authentication secret.
4. The authentication result is positive if the authenticating wireless node determines that the decrypted challenge text matches the challenge text originally sent in the second frame. The authenticating wireless node sends the authentication result.

Because the shared key authentication secret must be manually distributed and typed, this method of authentication does not scale appropriately in a large infrastructure network. Additionally, shared key authentication is not secure, because the shared key is stored in clear text, and thus it is not recommended for use.

WEP

Due to the nature of wireless LAN networks, securing physical access to the network is difficult. Unlike a wired network where a physical connection is required, anyone within range of a wireless AP can conceivably send and receive frames as well as listen for other frames being sent, making eavesdropping and remote sniffing of wireless LAN frames very easy. **Wired Equivalent Privacy (WEP)** is defined by the IEEE 802.11 standard and is intended to provide a level of data confidentiality that is equivalent to a wired network.

WEP provides data confidentiality services by *encrypting* the data sent between wireless nodes. WEP encryption for an 802.11 frame is indicated by setting a WEP flag in the MAC header of the 802.11 frame. WEP provides data integrity for random errors by including an **integrity check value (ICV)** in the encrypted portion of the wireless frame.

WEP defines two shared keys:

A multicast/global key	The multicast/global key is an encryption key that protects multicast and broadcast traffic between a wireless AP and all of its connected wireless clients.
A unicast session key	The unicast session key is an encryption key that protects unicast traffic between a wireless client and a wireless AP, and multicast and broadcast traffic sent by the wireless client to the wireless AP.

WEP encryption uses the RC4 symmetric stream cipher algorithm with 40-bit and 104-bit encryption keys. 104-bit encryption keys are not standard, however, although many wireless AP vendors support them.

To produce an encrypted frame, the following process is used:

1. A 24-bit integrity check value (ICV) is calculated that provides data integrity for the MAC frame.
2. The ICV is appended to the end of the frame data.
3. A 24-bit **initialization vector** (IV) is appended to the WEP encryption key.
4. The combination of [IV+WEP encryption key] is used as the input of a **pseudo-random number generator** (PRNG) to generate a bit sequence that is the same size as the combination of [data+ICV].
5. The PRNG bit sequence, also known as the key stream, is bit-wise exclusive ORed (XORed) with [data+ICV] to produce the encrypted portion of the payload that is sent between the wireless AP and the wireless client.
6. The IV is pre-pended to the encrypted [data+ICV] to create the payload for the wireless MAC frame. The result is IV+encrypted [data+ICV].

To decrypt the wireless MAC payload, the following process is used:

1. The IV is obtained from the front of the MAC payload.
2. The WEP encryption key is concatenated with the IV.
3. The concatenated WEP encryption key and IV is used as the input of a PRNG to generate a bit sequence of the same size as the combination of the data and the ICV (the same bit sequence as that of the sending wireless node).
4. The PRNG bit sequence is XORed with the encrypted [data+ICV] to decrypt the [data+ICV] portion of the payload.
5. The ICV for the data portion of the payload is calculated and compared with the value included in the incoming frame. If the values match, the data is considered to be valid (i.e., sent from the wireless client and unmodified in transit).

While the secret key remains constant over a long duration, the IV is changed periodically and as frequently as every frame. The periodicity at which IV values are changed depends on the degree of privacy required of the WEP algorithm. Changing the IV after each frame is the ideal method of maintaining the effectiveness of WEP.

802.1x is an IEEE standards-based framework for authenticating access to a network and, optionally, managing keys used to protect traffic. It relies on Remote Authentication Dial-In User Service (**RADIUS**), a network authentication service include in Windows Server 2003, to verify the network client's credentials with the domain controller and store the results in a database. The RADIUS server relies on the domain controller to authenticate the clients. 802.1x uses the Extensible Authentication Protocol (**EAP**) as a means of securing the conversation between the servers and clients and generating keys.

802.1x with **EAP-TLS** is a certificate-based system used to mutually authenticate wireless clients and RADIUS servers. It uses strong cryptographic keys to protect wireless traffic. This method requires public key certificates on the client and the RADIUS server. These public keys can be obtained from a trusted source third party or you can set up a Windows Server 2003 Certificate server to automatically enroll these certificates for clients using Active Directory.

802.1x with **PEAP (Protected Extensible Authentication Protocol)** can use Microsoft's Challenge Handshake Protocol version 2 (**MS-CHAPv2**) to provide secure password authentication without the use of certificates. This method works best in a small environment that does not have any certificate servers. It can also be used as an interim strategy to deploy a wireless network before implementing a certificate infrastructure.

Plan security for data transmission

Administrators spend a lot of time configuring secure access to data by using NTFS permissions, which does a great job of securing data from unauthorized users. When you transmit user data across the network, though, the data is wide open to a number of threats. Hackers can view that data using tools like Network Monitor and then use other tools to potentially alter it.

Let's first become familiar with common vulnerabilities that data is susceptible too and a view brief list of potential damage that can result.

Vulnerability	Damage
Data Modification	Modify data in transit
Denial of Service (DoS)	Floods a computer with TCP synchronization messages
Identity Spoofing	Send forged emails
Network Monitoring	Read clear text communication

Windows Server 2003 uses two encryption protocols for Virtual Private Networks (**VPNs**): Point-to-Point Tunneling Protocol (**PPTP**) and IP Security Protocol (**IPSec**). Since **L2TP** (Layer 2 Tunneling Protocol) does not have a built in data encryption, it relies on IPSec to accomplish this, while L2TP is left to set up the secure connection between the nodes.

The following table compares L2TP and PPTP:

L2TP	PPTP
Standards-based	Specific to Microsoft
Has header compression	Has no header compression
DES / 3DES compression	Built in, Microsoft proprietary compression
Supports Windows 2000, Linux, Solaris and others	Supports Windows
Requires only packet-based connectivity (this includes X.25, Frame Relay and ATM)	Requires IP-based internetwork
Uses IPSec encryption	Uses PPP encryption
Not fully compatible with NAT	Compatible with NAT

Secure data transmission between client computers to meet security requirements

Included in your plan for securing data communication in your network you must know what areas to protect. Plan on securing these layers:


- Application
 - ▶ SSL or TLS
 - ▶ SMB Signing
 - ▶ S/MIME
 - ▶ 802.1X
- Network
 - ▶ IPSec Protocol
 - ESP
 - AH
 - ▶ IPSec Mode
 - Tunnel
 - Transport

- ▶ Authentication protocols
 - Kerberos
 - Certificates
 - Preshared keys
- ▶ Encryption settings
- ▶ Filter lists
- ▶ Filter actions
 - Require security
 - Request security
 - Permit traffic
 - Block traffic
- Data Link and Physical
 - ▶ Replace hubs with switches to reduce broadcasts
 - ▶ Enable port authentication on switches
 - ▶ Restrict sensitive areas


Secure data transmission by using IPSec

IPSec works in two different modes, Transport and Tunnel mode. **Transport mode** is the default and should be used when you need to secure data transmission *within a network*. It will secure that traffic from endpoint to endpoint. **Tunnel mode** is used when you need to secure data transmission *between two networks*. It will secure the IP traffic from one network's border router to the next. This is very useful when you need to set up a secure connection for branch offices that may normally connect over the very unsecure Internet. Transport mode encapsulation and Tunnel mode encapsulation are depicted below.

Transport mode encapsulation:



Tunnel mode encapsulation:



Troubleshoot security for data transmission. Tools might include the IP Security Monitor MMC snap-in and the Resultant Set of Policy (RSoP) MMC snap-in

In Windows 2000 the **IP Security Monitor** ([ipsecmon.exe](#)) is a stand alone utility used to gather information about IPSEC. The tool shows active security associations, IPSEC Statistics and ISAKMP/Oakley Statistics. In Windows Server 2003 the IP Security Monitor has been replaced by the **IP Security Monitor** MMC. This new MMC has all the information contained in Windows 2000 but has a slightly different layout and more statistical information. Like many MMCs, the IP Security Monitor must be manually created by running **mmc.exe** and choosing File, then Add/Remove Snap-in.

For Windows 2000 and Windows XP clients, the **IPSec** command line configuration and scripting is done using a standalone tool called **IPSECCMD.EXE**. In Windows Server 2003 and Windows XP SP1, IPSec is part of the standard Network scripting commands shell called **Netsh**. From a command prompt, type **Netsh** to start netshell and then type **ipsec help** to see available commands and switches for the first level of IPSEC commands.

Resultant Set of Policies (RSoP) is a means of storing and viewing the combined effects of multiple GPOs (**Group Policy Objects**) on a single user or computer. The resultant settings are stored in the **WMI (Windows Management Instrumentation)** repository and viewable using the RSoP snap-in. The RSoP shows the resultant policy settings in the same hierarchical format as the GPO editor snap-in.

RSoP has two modes: **Logging** mode and **Planning** mode. Logging mode records the actual policy settings that have been applied to a user or computer. Planning mode allows scenarios to be run without actually applying GPOs to the target object. The latter mode allows “what-if” scenarios to be run. For example “what would be the RSoP for objects in a certain OU?” or “How would RSoP change by altering the security group filtering or WMI Query filtering?”

To launch the MMC snapin, click Start, Run, and type **rsop.msc**. Press Enter. The following dialog box is displayed as shown in *Figure 17*:

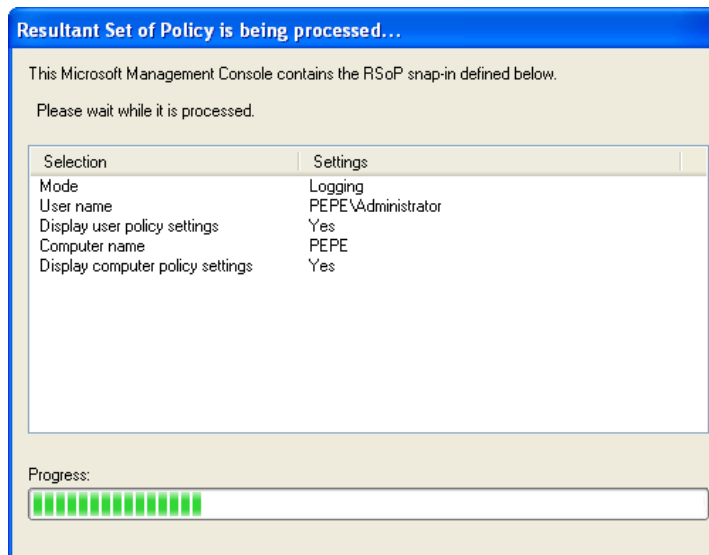


Figure 17 – Resultant Set of Policy

Another convenient method of gathering Group Policy settings is through a link in Help and Support. To gather the data in this way, click on the Tools link in Help and Support, followed by Advanced System Information, View Group Policy settings applied.

A progress bar is displayed in similar fashion as the MMC snapin. However the resulting output is now displayed in HTML format within Help and Support. At the bottom of this page is a link to save the output to an XML file. Using either method produces Group Policy Results in a easy to read and review format. *Figure 18* displays the Advanced System Information - Policy settings for PEPE.

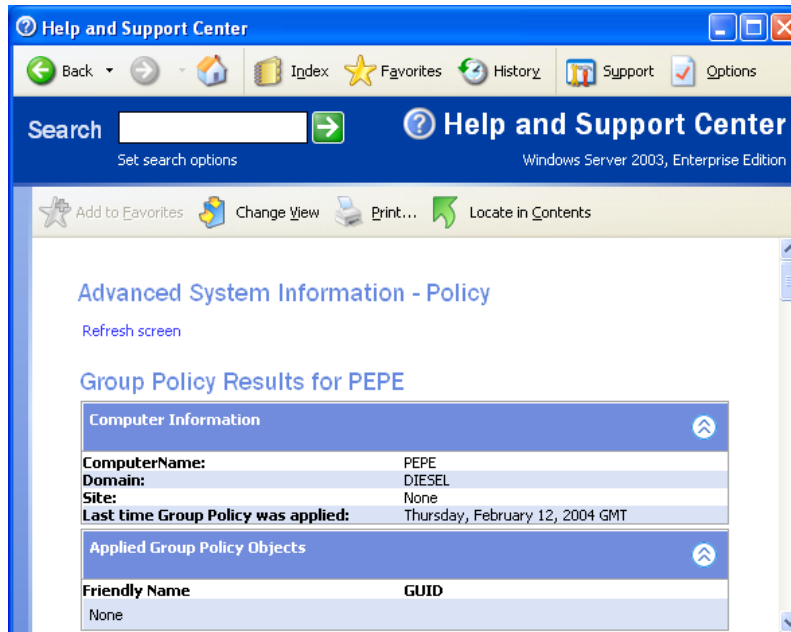


Figure 18 – Using Help and Support Center to View Group Policy Settings

Planning, Implementing, and Maintaining Security Infrastructure

Configure Active Directory service for certificate publication

If you decide that Active Directory is going to be used to publish certificates then you will need to know how Active Directory accomplishes that. Active Directory uses **DACLs** (directory access control lists) to either approve or deny a certificate request. For example, if a user requests a certificate for EFS, Active Directory will see if that user has the Enroll permission for that type of certificate. If the user does, then her request is approved. This happens quickly, and does not rely on an administrator to manually check to see if that user had the appropriate permission to enroll.

You can also set up Group Policy to automatically enroll User and Computer certificates in Windows Server 2003. Previously, in Windows 2000 Server, you were only allowed to auto-enroll Computer certificates. *Figure 19* shows the Autoenrollment Property settings.

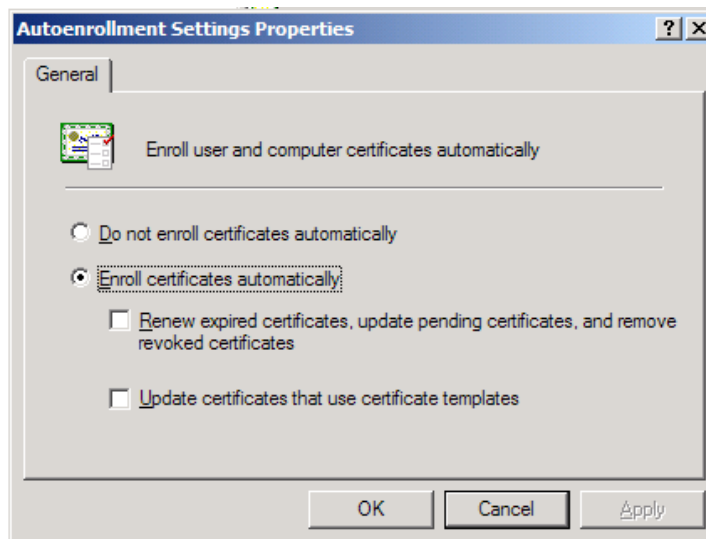


Figure 19 – Enrolling User and Computer Certificates

Plan a public key infrastructure (PKI) that uses Certificate Services

The Windows Server 2003 Public Key Infrastructure enables you to deploy strong security solutions that use digital certificates and public key technology. Security solutions can include the following:

- Secure mail, which uses certificates and the **Secure/Multipurpose Internet Mail Extensions (S/MIME)** protocol to ensure the integrity, origin, and confidentiality of e-mail messages.
- Secure Web sites, which use certificates and certificate mapping to map certificates to network user accounts for controlling user rights and permissions for Web resources.
- Secure Web communications, which use certificates and the **Secure Sockets Layer (SSL)** and **Transport Layer Security (TLS)** protocols to authenticate servers, to optionally authenticate clients, and to provide confidential communications between servers and clients.
- Software code signing, which uses certificates and digital signing technology (such as Microsoft Authenticode) to ensure the integrity and authorship of software that is developed for distribution on an intranet or on the Internet.
- Smart card logon process, which uses certificates and private keys stored on smart cards to authenticate local and remote access network users.
- Internet Protocol security (**IPSec**) client authentication, which has the option to use certificates to authenticate clients for IPSec communications.
- Encrypting File System (**EFS**), which uses certificates for both EFS user and EFS recovery agent operations.
- Custom security solutions, which use certificates to provide confidentiality, integrity, authentication, or nonrepudiation.

PKI provides security for conversations between computers. PKI relies on these components to achieve its purpose:

- **Digital signature (certificate)** – Like a person's cursive signature, this verifies that a message is actually from the stated source. In a networking computer environment, users, computers, routers and organizations can use digital certificates to certify their identities. Certificates thus provide **authentication**.
- **Encryption** – Like the snapping lid on a bottle of Ketchup verifies that the tasty red contents have not been tampered with, encryption verifies that the message has not been corrupted or viewed. Encryption thus provides **privacy** as well as **integrity** to the communicating entities.

Identify the appropriate type of certificate authority to support certificate issuance requirements

A **Certificate Authority (CA)** issues certificates. In addition to trusted third-party commercial companies such as [VeriSign](#), a Windows Server 2003 system can also be configured to issue certificates and to verify that an existing certificate is legitimate and belongs to the entity it says it is from. Before issuing a certificate, a CA must validate the applicant's identity.

Here are some preliminary items to review before installing a CA in Windows Server 2003:

- Reconsider installing a CA on a Domain Controller, as this could result in an overly busy system.
- Ensure you're happy with the computer name: it can't be renamed after loading Certificate Services nor can the computer join or leave a domain.
- Have a unique CA name handy for each CA in your enterprise.

There are two main types of Windows Server 2003 CAs, each of which can be a **Root** or **Subordinate Certificate Server**:

- **Enterprise CA**
 - ▶ Integrates with Active Directory – Certificates and **Certificate Revocation Lists (CRLs)** are published in AD. Moreover, certificates can be issued only to objects in the Active Directory forest. Use a Standalone CA for objects outside of the forest
 - ▶ Install an Enterprise Root CA before all other CAs since they rely on the Root CA to certify them. Also, the server that is to hold the Enterprise Root CA needs to have its computer account put in the **Cert Publishers group**. The installation has to be done by an administrator in the **Enterprise Admin group**
 - ▶ Automatically approves certificates: based on user account and group account information and certificate template information
 - ▶ Works with **smart cards**
 - ▶ Should have some fault tolerance built in, such as regularly scheduled backups
 - ▶ Its associated server name becomes part of the certificates it manages, thus you can't change the server name after installing Certificate Services
 - ▶ Doesn't require the person requesting a certificate to supply all identifying information since this information is apparent based on the user's logon account. Moreover, the certificate type is based on the **certificate template**
- **Standalone CA**
 - ▶ Doesn't require Active Directory
 - ▶ Supports **S/MIME** (secure email), Secure Sockets Layer (**SSL**) and Transport Layer Security (**TLS**)
 - ▶ Works with external networks (**extranets**)
 - ▶ Holds requests for certificates in a **Pending Queue** for later (manual) approval by a CA

administrator and thus does not immediately grant certificates

- ▶ Requires a person requesting a certificate to supply all identifying information as well as the type of desired certificate
- ▶ Doesn't use certificate templates

Plan the enrollment and distribution of certificates

Before a CA issues a certificate, a request must first be made. In some cases, this process is automatic (e.g., a smart card logon to a domain), while in other cases user interaction is required.

Run the **Certificates** MMC and select **Certificates** -> **Personal** -> **All Tasks** -> **Request New Certificate**. You are then presented with a **certificate template** (i.e. policy) choice.

These templates control the issuing of a certificate and may include additional choices in the dialog box if you added additional functionality, such as a smart card. Lastly, you assign a friendly name and description. If the certificate is granted to a user, the user can **cancel**, **install** or **view** the certificate. This method of requesting a certificate works only with an Enterprise CA.

The other manual way is via the **Certificate Services Web Page**, which is accessible through **http://<servername>/certsrv**, where **servername** is the name of the server hosting Certificate Services as shown in the Advanced Certificate Request dialog box in *Figure 20*. From this page, you can choose to **Request a certificate**. This works with a Standalone or Enterprise CA.

An **Enterprise CA** will either grant or refuse to grant a certificate request. In the case of a grant, the requesting user is asked to install the certificate. On the other hand, a **Standalone CA** will place the request in a pending state so that an administrator can later deal with it. One advantage of requesting a certificate from an Enterprise CA is that the user gets an automatic response as to the success of the request. To check for a pending certificate, type **http://<servername>/certsrv** and then select **Check on a pending certificate**. You will be presented with a list of certificates in one of these states: denied, issued or pending.

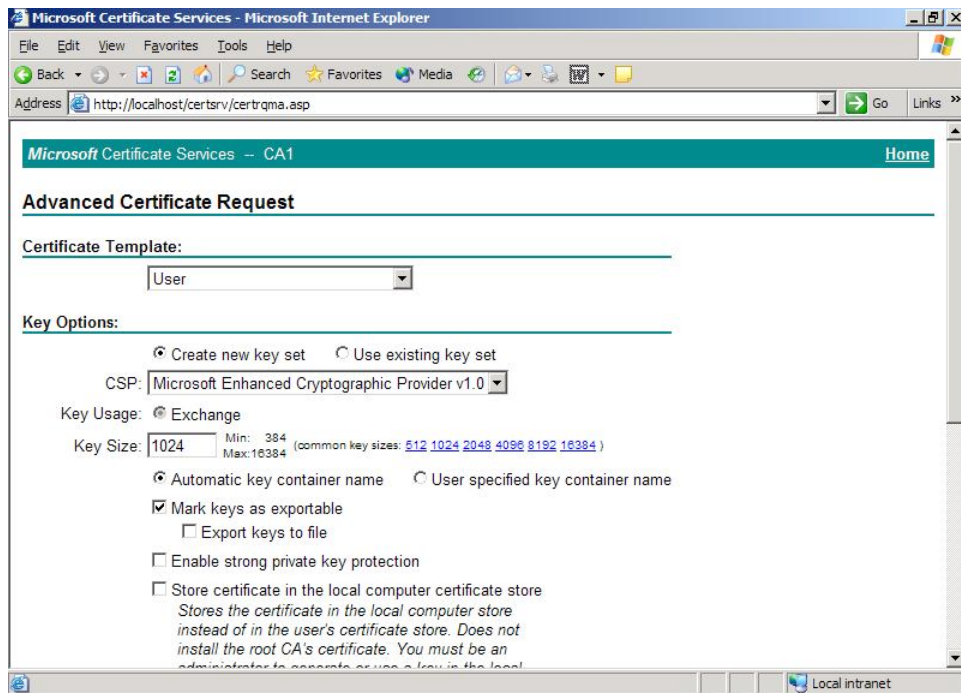


Figure 20 – Microsoft Certificate Services

Plan for the use of smart cards for authentication

A smart card is a tamper-resistant device that includes a built-in microprocessor, operating system, and memory that is used to securely store personal information.

Smart card logon is a *two-step* authentication mechanism that uses a hardware device, known as smart card, to store a user's public key credentials, and a **Personal Identification Number (PIN)** as the secret key to authenticate the user to the smart card. A smart card offers tamper resistant storage space for the user's private key that can only be accessed by entering a secret Personal Identification Number (PIN), and cryptographic engine support for operations such as digital signatures and key exchange. A PIN is only used to authenticate the user to the smart card and it is never sent over the network, as is the case of shared secret logon where a derivative of the user's password is transmitted over the network. Once authenticated, the user's public key credentials are retrieved from the card by the Windows Security Subsystem and verified to make sure it is valid and that a trusted party issued it.

A user attempting to logon to a Windows Server 2003 domain using a smart card, must go through the following sequence of steps before they can successfully be logged on.

1. User inserts smart card into reader.
2. The WINLOGON service traps the smart card arrival event and dispatches it to MSGINA (Microsoft Graphical Identification and Authentication).
3. MSGINA prompts the user to enter his personal identification number (PIN).
4. The user types in his PIN.

5. MSGINA sends the user-supplied PIN to the Local Security Authority (LSA).
6. LSA uses the PIN to access the smart card and retrieve the user's certificate.
7. Once the user's public key credentials are retrieved, the Kerberos **Security Service Provider (SSP)** on the local machine sends a signed user certificate to the **Key Distribution Center (KDC)**.
8. The KDC compares the user's certificate with the certificate that is stored in Active Directory. Since the user certificate is signed with the user's private key, the KDC can validate the integrity of the client certificate.
9. Once the user certificate is validated, the KDC generates a logon session key, encrypts it along with the **TGT** (ticket granting ticket), using the public key extracted from the client certificate, and sends the encrypted contents to the client. This guarantees that only the holder of the private key can decrypt the logon session key.
10. The client receives the encrypted logon session key and TGT and uses its private key to decrypt them. Once decrypted the client can present the TGT to the **TGS** (ticket granting service). Now that the client is in possession of the logon session key, all Kerberos communication will now use symmetric encryption.

Plan a framework for planning and implementing security

Microsoft has made it easy to plan and implement security. They've done so by making their best practices available to us through what is known as **Microsoft Solutions Framework (MSF)**. The MSF consists of people, processes, and risk management. Each part of the framework plays a key role in the overall design. It is essential that all involved people can communicate effectively. You should involve a representative sample of management and other users to pull information from many sources in your organization. These people should be able to assist in aligning technology solutions with business requirements.

MSF is a suite of guidelines and principles that provide models to build and deploy a distributed network. It consists of three phases that can each be repeated many times over the life of a network. The three phases of the MSF framework are as follows:

The Planning Phase

In the planning phase, you need to assess your current security model and policies and then decide what changes should be made to fulfill the needs of the organization. This is the time for all key personnel involved to ask for what they need regardless of how it will be provided. The goal is to create a vision for optimum security for the network. Each person should understand the overall goal and the components to reach the goal. You should establish a system of measurable metrics to help you stay on track toward the goal. You might want to meet in a conference room with a white board or flipcharts and any relevant documentation. All ideas should be considered, no matter how far-fetched they might seem at first.

The Building Phase

As you can imagine, the personnel responsible for building the security design are not generally the same people that were involved in the planning. In the building process, experienced network administrators, security specialists, and consultants will use the hardware and software at their disposal to attempt to create the security that the security design team has envisioned.

These administrators will create, deploy, and test security templates and policies. This phase is usually performed in a testing lab to assess the impact of security decisions before placing them into a production environment. This is done to protect the productivity of the workers. After successful completion in a lab, the templates should be rolled out in phases. You should use a test group to make sure that you haven't missed anything in the lab before rolling the security out to the whole network. Knowledge of what is available and how it relates to network security is the key component in the building phase.

The Managing Phase

After implementation of the security design, you are responsible for managing the design to make sure that it provides the security that was envisioned by the security design team. If the network is never going to grow or change again, this would be relatively simple. However, because the network is constantly changing and growing, you should make sure that each change is carefully considered in regard to its impact on your new security design. As mentioned previously, a flaw in one area of the network can have an effect on the entire network.

The help desk will generally provide end-user support, but you are responsible for monitoring the network to make sure that the new policies are being enforced. You will likely detect security vulnerabilities and use the system's features to better protect the network. You should inform the appropriate managers of any misuse of the network to ensure that the policies are enforced. Success in this phase depends on constant monitoring and taking personal responsibility for the security of the network.

Plan for security monitoring

When you monitor any network for security the key is to know what to look for. You will want to become familiar with patterns that emerge and also monitor for well-known types of threats. The list below is a list of common types of threats; there are other types of threats out there.

Types of Threats	Examples
Spoofing	Forge email addresses and replay authentication packets
Tampering	Alter data during transmission and change data in files
Repudiation	Delete a critical file and deny doing it; purchase a product and later deny doing it
Information disclosure	Expose information in messages; expose code on Web sites
Denial of Service (DoS)	Flood network with SYN packets; flood network with forged Internet Control Message Protocol (ICMP) packets (pings)
Elevation of privilege	Exploit buffer overruns to gain system privileges; exploit the Local System account

Plan a change and configuration management framework for security

A planned change and configuration management framework for security will always anticipate and meet management goals and end user needs. A plan should include the following features:

- User Data Management
- User Settings Management
- Software Installation and Maintenance
- Patch Management
- End User Training
- Security Settings
- OS Installation guides

Plan a security update infrastructure. Tools might include Microsoft Baseline Security Analyzer and Microsoft Software Update Services

You can use **Microsoft Baseline Security Analyzer (MBSA)** to scan for security-related updates on multiple computers. MBSA Version 1.1.1 includes both a GUI tool and a command-line interface tool. You can use these tools to perform scans of Windows systems on your network. MBSA runs on Windows 2000, Windows XP, and Windows Server 2003 systems. You can perform scans of all Windows NT-based clients including Windows NT Workstation and all later clients. You can also scan for updates to applications running on the clients including Internet Explorer (IE), and Office applications including Office 2000 and later.

Software Update Services (SUS) is new to Windows Server 2003 but backward compatible to Windows 2000 servers running Service Pack 3 or later. You can use SUS to update clients running Windows 2000 Professional and Windows XP Professional with the latest service packs, system fixes and security patches. SUS enables an administrator to automatically download, test, approve, and install the latest critical updates and service packs from the Microsoft Windows Update Web site. You should be familiar with the features of SUS, as shown in the table below.

SUS Feature	Brief Description
Built-in security	The administrative pages of SUS are Web based through IIS and restricted to local administrators on the computer that hosts the updates. The synchronization always validates the digital certificates on any downloads to the update server. Any files that are not from Microsoft are automatically deleted.
Selective content approval	Administrators can selectively choose which updates they want to be available to clients. This allows you to test the packages before deploying them.
Content synchronization	You receive the latest critical updates and service packs from Microsoft through the process of synchronization. You can set a schedule for automatic synchronization at preset times. Alternatively, you can use the Synchronize Now button to manually synchronize the server.
Server-to-server synchronization	You can point your server to another server running Microsoft SUS instead of to the Windows update server. This creates a single point of entry for updates into the network, without requiring that each SUS server download updates from the external Microsoft source.
Multi-language support	SUS supports the publishing of updates to multiple operating-system language versions.
Remote administration via HTTP or HTTPS	The SUS administrative interface is Web based. This allows you to manage it remotely as if you were sitting in front of the server itself. Remote administration requires Internet Explorer 5.5 or later.
Update status logging	You can specify the address of a Web server where the Automatic Updates client should send statistics about updates that have been downloaded and installed.

Practice Questions

Chapter 1 Planning and Implementing Server Roles and Server Security

1. You are a systems administrator for your company. You have been tasked with comparing your standard desktop build currently in use on the desktop systems in your enterprise against the default configuration security settings that are applied during a new installation of the Windows XP Professional operating system. What is the easiest way to accomplish this task by using a script? Select the best answer.

- A. Use the Security Configuration and Analysis tool against all of the systems.
- B. Use MBSA against all of the systems.
- C. Use SIGVERIF.exe against all of the systems.
- D. Use SFC.exe against all of the systems.
- E. Use Secedit against all of the systems.

2. You are a domain administrator for an airline company and you have been asked to review the design of your internal DNS servers and the type of zone updates that are in use. Currently the DNS zones are configured across your DNS servers in a standard primary and standard secondary configuration.

All of the group memberships in the domain have been configured with the default settings and with default memberships. You need to ensure that dynamic updates are allowed for all of the zones on the internal DNS servers.

Which of the following steps allow you to accomplish your task using the least amount of administrative effort? You need to perform your task via the command line using security best practices and a user account with the minimum level of required administrative privilege to perform the task.

Select all that apply.

- A. Using the RUNAS service, go to Start, click Run, and type `runas /user:chicklittle\<ACCT> cmd` (where <ACCT> is a member of the DnsAdmins group on the DNS server).
- B. Using the RUNAS service, go to Start, click Run, and type `runas /user:chicklittle\<ACCT> cmd` (where <ACCT> is a member of the Server Operators group on the DNS server).
- C. Using the RUNAS service, go to Start, click Run, and type `runas /user:chicklittle\<ACCT> cmd` (where <ACCT> is a member of the Domain Administrators group on the DNS server).
- D. Change the zone type of your DNS zones to Active Directory integrated.
- E. From the command line, enter `dnscmd <SERVERNAME> /Config <ZONENAME> /AllowUpdate 1`
- F. From the command line, enter `dnscmd <SERVERNAME> /Config <ZONENAME> /AllowUpdate 2`

3. You are a systems administrator for your company and have been asked to compare your standard desktop build currently in use on the desktop systems in your enterprise against the default configuration security settings that are applied during a new installation of the Windows XP Professional operating system.

You need to outline for management the default local security policy settings for auditing on a new installation of the Windows XP Professional. What are the proper settings for a new installation of the Windows XP Professional?

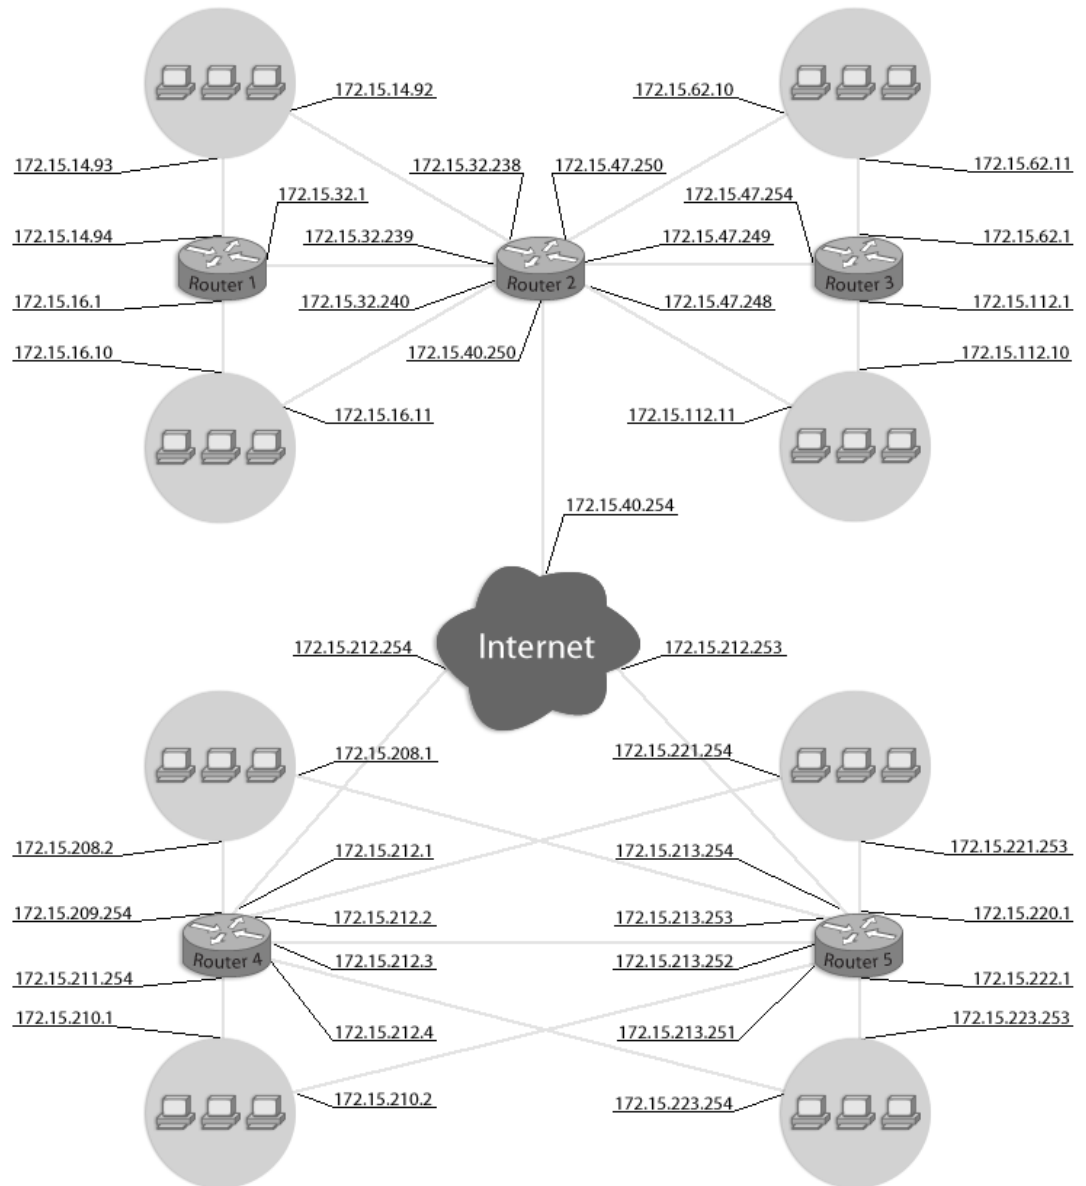
Select all that apply.

- A. Audit account logon events -Success
 - B. Audit account logon events -Success and Failure
 - C. Audit account logon events -No Auditing
 - D. Audit logon events -Success
 - E. Audit account management -No Auditing
 - F. Audit account management -Failure
4. You are the systems administrator for your company. Recently, a number of changes have been made to the base configuration of a Windows Server 2003 production file server that has been installed in the SYSTEST workgroup. The system has a total of three partitions. The C:\ partition (where the system files were installed) and the D:\ partition (where the boot partition was configured) were both formatted as NTFS. The E:\ drive was originally formatted with FAT32 and has numerous programs installed and configured on it currently. To increase the level of effective security on the system you have run the CONVERT utility on the E:\ partition as follows: convert E:/FS:NTFS /V What additional actions, if any, would you need to take so that the default file system ACLs are properly configured on this converted partition with the least amount of administrative effort?
- Select the best answer.
- A. Use the Security Configuration and Analysis MMC snap-in to import the defltsv.inf security template to the local policy.
 - B. Use the Local Computer Policy MMC snap-in to import the defltsv.inf security template to the local policy.
 - C. Use the Local Computer Policy MMC to import the Setup security.inf security template to the local policy template.
 - D. Import the Basicsv.inf template into a GPO and apply it to the domain object in the Active Directory.

Chapter 2 Planning, Implementing, and Maintaining a Network Infrastructure

1. You are the network administrator for a company whose network infrastructure is laid out in the exhibit. You have been asked to set up the new remote location that has been configured using the 172.15.208.0/20 IP address block.
- The company headquarters network configuration shown at the top of the diagram is connected to the Internet and the remote office via a Windows Server 2003 Routing and Remote Access Server (not shown) at IP address 172.15.40.254.
- The company's remote office is shown at the bottom of the diagram and is connected to the Internet and the company's headquarters via a multihomed Windows Server 2003 Routing and Remote Access Server (not shown) at IP addresses 172.15.212.253 and 172.15.212.254.
- You need to review and resolve any single points of failure other than the failure of the single Routing and Remote Access Server itself on the current network.
- How would you alleviate the single points of failure on this network, if any, other than the hardware that makes up the Routing and Remote Access Server, which might isolate subnets from reaching other subnets or the Internet?
- Select the best answer.
- A. Add an additional router at the remote location between routers R4 and R5.
 - B. Add two additional routers at company headquarters: one between routers R1 and R2 and another one between R2 and R3.
 - C. Add an additional network segment from R2 and the Windows Server 2003 Routing and Remote Access Server at IP address 172.15.40.254.
 - D. Add an additional Internet network connection point at IP address 172.15.40.254.
 - E. Nothing further needs to be done because there are no single points of failure present.

Exhibit(s):



2. You are a desktop administrator for gunderville.com. Your organization is made up of three sites; two are remote offices and one is the company headquarters. The two remote locations are made up of four subnetworks each; interconnected internally by layer 2 switches and connected to one another and headquarters by routers. The main company headquarters consists of eight subnets interconnected internally by layer 3 switches.

The client systems on your network are running a number of operating systems including Windows 98, 2000, XP, and Server 2003. There are two WINS servers and two DNS servers local to each subnet and the clients are always configured to use those servers for name resolution only. Recently, new server hardware was installed for SRV007 to replace a legacy system and there have been reported issues with incoming connectivity to the server. The server itself seems to be able to initiate connections to systems fine. You attempt to PING SRV007.gunderville.com and you receive the following error:

Pinging SRV007.gunderville.com [149.88.72.19] with 32 bytes of data:

Request timed out.

Request timed out.

Request timed out.

Request timed out.

Ping statistics for PING SRV007.gunderville.com:

Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)

You then log in to SRV007

gunderville.com and run ipconfig and review the output

Ethernet adapter Wireless Connection:

Connection-specific DNS Suffix .gunderville.com

IP Address.: 149.88.72.91

Subnet Mask.: 255.255.255.0

Default Gateway: 149.88.72.1

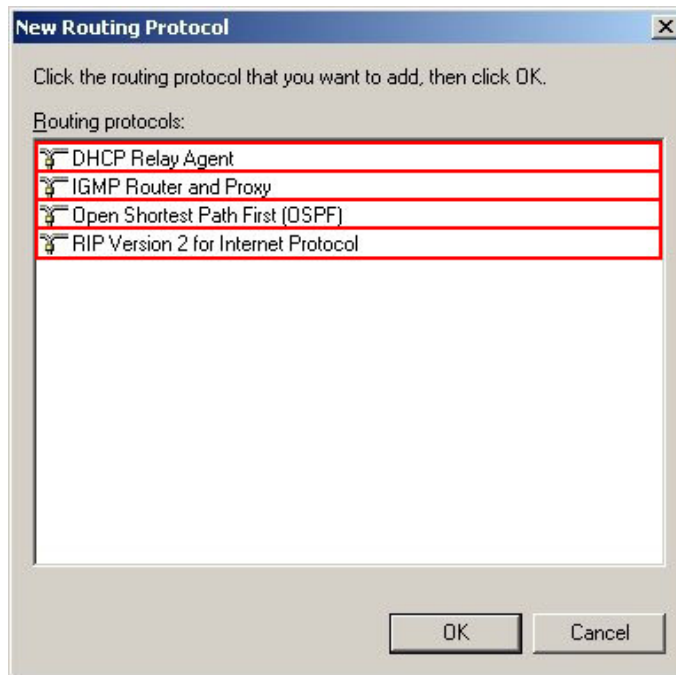
What are the most likely reasons that this system cannot be connected to but the system itself can make outbound initiating connections?

Each answer presents a possible solution. Choose two.

- A. The default gateway is incorrect.
- B. The server's IP address is not assigned correctly.
- C. Routing for the network is inconsistent or beginning to fail.
- D. The IP address entry for the DNS server is incorrect.
- E. The subnet mask is invalid for this subnet

3. You are a desktop administrator for gunderville.com with client systems running a number of different operating systems, including Windows 98, 2000, XP, and Server 2003.
- There are two WINS servers and two DNS servers local to each subnet and the clients are always configured to use those servers for name resolution only.
- Recently a new server was installed (SYS03) on the 18.12.16.0/21 network but there have been issues with network communications from this system; all of the other communications on this subnet are occurring normally.
- You decide to check the IP configuration of SYS03 by running IPCONFIG/ALL and receive the following information:
- ```
Connection-specific DNS Suffix . : gunderville.com
Description : Intel(R) PRO/100 VE Network Connection
Physical Address. : 05-10-H9-C1-A2-AB
Dhcp Enabled. : No
Autoconfiguration Enabled : No
IP Address. : 18.12.23.253
Subnet Mask : 255.255.248.0
Default Gateway : 18.12.24.24
DHCP Server :
DNS servers : 18.12.22.253
18.12.21.252
Primary WINS server : 18.12.20.252
Secondary WINS server : 18.12.19.252
Lease Obtained. : Monday, September 29, 2003 12:50:53 PM
Lease Expires : Tuesday, September 30, 2003 12:50:53 PM
```
- What is the most likely issue with the IP address configuration for this system?  
Select the best answer.
- A. The server's IP address is not correctly assigned for the subnet.
  - B. The default gateway is incorrect.
  - C. Routing for the network is inconsistent or beginning to fail.
  - D. The IP addresses entries for the DNS servers are incorrect.

4. You are a network administrator for gunderville.com. You are designing the network configuration for your company at a new subsidiary. You have used the 152.166.25.0 IP address that your company owns and segmented the networks in this new location to allow for 450 hosts per subnet (current) and also allow for an anticipated growth of 40% for the client systems.
- The overall network design of this location and to other remote locations is set up for 17 hops between the two most remote subnets. You need to configure a routing solution for your design that will automatically allow for the update of routers in use.
- Which update method for routing can be used to dynamically update the routing information for your network so that static routing entries do not have to be manually maintained?
- To answer, select the appropriate routing protocol in the exhibit.



5. You are a network administrator for your Windows Server 2003 domain running at Windows 2000 Server domain functional level (native mode). You are designing the network configuration for your company at a new subsidiary.

You need to use the 152.166.25.0 IP address that your company owns and segment the networks in this new location in such a way that allows for 450 hosts per subnet currently and also allows for an anticipated growth of 40% for the client systems.

Although the number of clients is expected to grow, the number of required subnets at this location is not expected to change—no more than 30 subnets will ever be needed.

You need to adjust for the anticipated growth numbers for clients. You should not calculate for any additional subnets than are necessary for the network configuration or for the total number of hosts per subnet whenever possible, yielding on the side of the host per subnet in the event of a configuration conflict. Which subnet mask length and subnet mask should be used for the subnets at this location?

To answer, drag the correct subnet mask length and subnet mask to the appropriate fields in the exhibit.

A.255 B.254 C.0 D.252 E.248 F.23 G.21 H.22

**New Scope Wizard**

**IP Address Range**  
You define the scope address range by identifying a set of consecutive IP addresses.

Enter the range of addresses that the scope distributes.

Start IP address:

End IP address:

A subnet mask defines how many bits of an IP address to use for the network/subnet IDs and how many bits to use for the host ID. You can specify the subnet mask by length or as an IP address.

Length:

Subnet mask:

< Back   Next >   Cancel



## Chapter 3 Planning, Implementing, and Maintaining Routing and Remote Access

1. You are a network administrator for your Windows Server 2003 domain and you are updating the network configuration for your company. Part of your design calls for outlining the appropriate level of security of data that is transmitted on the network and adding that information into the standards documentation so that developers can properly reference this information as they create their applications that will be required to run under this configuration.

You need to determine which types of authentication protocols can be used when you need to ensure that a one-way encrypted password, mutual authentication process is used where the authenticator sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

Which authentication protocol meets all of these specifications and can be used in this situation?

Select the best answer.

- A. Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v1) version 1
  - B. Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2) version 2
  - C. Protected Extensible Authentication Protocol (PEAP)
  - D. Extensible Authentication Protocol (EAP)
  - E. Challenge Handshake Authentication Protocol (CHAP)
2. You are a network administrator for your Windows Server 2003 domain and you are updating the network configuration for your company. You have used part of the 191.99.74.0/16 IP address range that your company owns and segmented the networks in this new location in such a way to allow for 60 hosts per subnet and also allow for an anticipated growth of 15% for the client systems.

The overall network design of this location and to other remote locations is set up so that the distance between the two most remote subnets is 13 hops between all 44 routers in use on the LAN. You need to configure a distance vector routing solution for your design that will automatically allow for the update of the routing information to all routers that are in use, as well as provide authentication and multicast announcing support.

Which of the following routing protocols can be used to dynamically update the routing information for your network so that static routing entries do not have to be manually maintained and meet all of the listed requirements?

Select the best answer.

- A. Routing Information Protocol (RIP) version 1
- B. Routing Information Protocol (RIP) version 2
- C. Open Shortest Path First (OSPF)
- D. Interior Gateway Protocol (IGP)

3. You are a network administrator for your Windows Server 2003 domain and you are updating the network configuration for your company. You have used the 190.25.77.0 IP address that your company owns and segmented the networks in this new location in such a way so as to allow for 60 hosts per subnet currently and also allow for an anticipated growth of 15% for the client systems. Your design also calls for adding to your design a routing configuration that will allow you to forward IP multicast traffic.

You have decided to use the Routing and Remote Access Service on your Windows Server 2003 system to handle this requirement in your environment. As part of the design and configuration of this routing solution, you need to map out the TTL thresholds for the multicast scopes that will be used. TTL boundaries prevent the forwarding of IP multicast traffic with a TTL less than a specified value and they apply to all multicast packets regardless of the multicast group. What is the TTL threshold for a multicast restriction to a single site?

Select the correct answer.

- A. 15
  - B. 1
  - C. 63
  - D. 127
  - E. 191
  - F. 255
4. You are a network administrator for your Windows Server 2003 domain and you are updating the network configuration for your company. You have used the 190.25.77.0 IP address that your company owns and segmented the networks in this new location in such a way so as to allow for 60 hosts per subnet currently and also allow for an anticipated growth of 15% for the client systems. Your design also calls for adding to your design a routing configuration that allows you to forward IP multicast traffic and use the Internet Group Management Protocol (IGMP).

You have decided to use the Routing and Remote Access Service on your Windows Server 2003 system to handle this requirement in your environment. What are the proper steps for enabling IGMP for use?

To answer, drag the appropriate steps to the right in the proper order.

- A. Open the Configure your Server Wizard. \_\_\_\_\_
- B. Open the Routing and Remote Access MMC snap-in. \_\_\_\_\_
- C. Select General under Remote Access Policies and then right-click and choose New Routing Protocol. \_\_\_\_\_
- D. Select General under IP Routing, right-click, and choose New Routing Protocol. \_\_\_\_\_
- E. In the Select Routing Protocol dialog box, click Add, New Protocol, choose IGMP Router and Proxy, and then click OK.

## Chapter 4 Planning, Implementing, and Maintaining Server Availability

1. You are the server administrator for TERMSRV01 which is a Windows Server 2003 Terminal Server. You have decided that you need to perform a baseline reading of the system with System Monitor in an effort to have the baseline available in the future for comparison if needed and you have decided that system memory is going to be the biggest point of contingency in the near future.

Which System Monitor counters do you need to enable to get information on the overall rate at which the hard disk is used (frequency of disk access)?

Choose four.

- A. Processor\% Processor Time
- B. Network Interface Connection\Bytes Total/sec
- C. Memory\Page Reads/sec
- D. Memory\Pages Input/sec
- E. PhysicalDisk\% Disk Time
- F. PhysicalDisk\Current Disk Queue Length

2. You are the domain administrator for your Windows Server 2003 domain and have been asked to design and deploy a Windows clustering solution for your Terminal Services configuration. Your configuration calls for the ability to extend cluster shared disk partitions on your configuration as needed without disrupting the current data on the volumes. Which native tool allows you to do this?

Select the best answer.

- A. MSMQ
- B. COMPMGMT.MSC
- C. DISKMGMT.MSC
- D. DISKPART.EXE

3. You are the server administrator for TSSYS04, which is a Windows Server 2003 Enterprise Edition Terminal Server. Your backup strategy utilizes one full backup and then daily incremental backups. TSSYS04 has eight 36GB SCSI3 10,000RPM drives; five of the drives make up the data array called array 2. Disk3, disk4, disk5, disk6, and disk7 are all used to store data on its own partition and are part of a RAID 5 configuration set up through the Windows Server 2003 operating system. The system and boot partitions are on two of the other disks (disk0 and disk1) and are part of a RAID 1 implementation that is driven by the server hardware controller. This is array 1 on the system. Disk2 is a hardware-driven, online spare for the mirrored system and boot drives that make up the first array. The rebuild of data to the online spare is calculated at a maximum of 75 minutes in the event of a failure.
- The full back up is performed each Saturday at 9:00PM EST. The full back up takes 2 hours. The daily incremental backups are performed each day at 9:00PM, Monday through Friday. There is no full backup nor incremental backup performed on Sunday. The amount of time needed for the incremental backups varies but it never exceeds 75 minutes.
- TSSYS04 suffers a power spike at approximately 6:00PM on a Saturday and has immediate hardware issues. It suffers a hard disk failure at 6:05PM on disk1 at approximately 6:00PM and at 6:20PM disk3 also fails. At 7:20PM disk2 also fails. What needs to be done to bring the server back to working order so that users can access the data on the server?

Select the best answer.

- A. The failed drives need to be replaced and the data needs to be restored from the full backup.
  - B. The failed drives need to be replaced and the data needs to be restored from the full backup and the last incremental backup.
  - C. The failed drives need to be replaced and the data needs to be restored from the full backup and all of the incremental backups.
  - D. The failed drives will need to be replaced. Nothing additional needs to be done because the RAID configuration allows the system to continue to run.
4. You are the domain administrator for a Windows Server 2003 domain and have been asked to design and deploy a Windows Clustering solution for your Terminal Services configuration. Your current configuration is using a total four Windows Server 2003 Enterprise Edition servers with four Processors in each system and 8GB of RAM. You need to increase this cluster solution to its maximum possible threshold for both server hardware and cluster size.
- What are the maximum settings with regard to the number of clustered systems, processor, and memory hardware under the Windows Server 2003 Enterprise Edition operating system?
- Select the best answer.
- A. The maximum configuration of the cluster would be a total of eight cluster servers with each cluster server having a maximum total of eight CPUs and 32GB of RAM.
  - B. The maximum configuration of the cluster would be a total of four cluster servers with each cluster server having a maximum total of eight CPUs and 32GB of RAM.
  - C. The maximum configuration of the cluster would be a total of four cluster servers with each cluster server having a maximum total of four CPUs and 32GB of RAM.
  - D. The maximum configuration of the cluster would be a total of four cluster servers with each cluster server having a maximum total of eight CPUs and 8GB of RAM.

## Chapter 5 Planning and Maintaining Network Security

1. You are the domain administrator for gunderville.com. Client systems in use include Windows 2000 Professional and Windows XP Professional for most locations. For branch office one, the exception would be six Windows 98 clients that are still in use. Server systems in use include 2000 Server and Server 2003.

Branch office one has all of the clients using DHCP locally and interconnected locally by a layer 2 switch and back to the main office via a router and a private leased line. There are a total of 43 host systems at this location including the installed servers.

Branch office two has all of the clients using DHCP and connected locally by a layer 3 switch and back to the main office via a Windows Server 2003 Routing and Remote Access server and a ISP connection to a VPN server at company headquarters. There is a slower, secondary demand-dial connection back to the main office that is used if the primary connection goes down. There are a total of 39 host systems at this location including the installed servers.

Branch office three has all of the clients using DHCP and connected locally by a layer 3 switch and back to the main office via an ISP connection to a VPN server at company headquarters. There are a total of 22 host systems at this location including the installed servers.

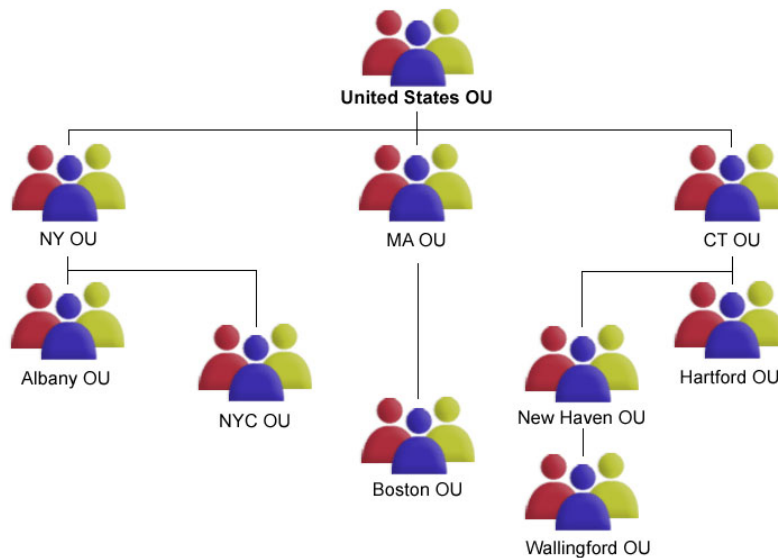
You need to provide a security configuration for data transmissions that occur over the Internet to the main company headquarters so that all of this data is secured against tampering and unauthorized modification.

How would this be accomplished in your environment by performing only the required steps, not adding any additional resource load to the servers or networks than is absolutely necessary to accomplish the required tasks, and using the least amount of administrative effort?

Select all that apply.

- A. The Microsoft L2TP/IPSec VPN Client needs to be installed on all Windows 98 systems using the Dial-up Networking version 1.4 upgrade if they were making the connections to the VPN server individually.
- B. The Microsoft L2TP/IPSec VPN Client needs to be installed on all Windows 98 systems making VPN connections to the company headquarters via the RRAS server connection.
- C. The Microsoft L2TP/IPSec VPN Client would not be needed for any of the systems in use.
- D. The Microsoft L2TP/IPSec VPN Client needs to be installed on all Windows 98 systems making the connections to the VPN server individually by way of the PPTP adapter.
- E. For securing the required traffic, the Encapsulating Security Payload (ESP) is needed as the security method.
- F. For securing the required traffic, the Authentication Header (AH) is needed as the security method.

2. You are the enterprise administrator for gunderville.com. Client computers run Windows XP Professional and Windows 2000 Professional. The Active Directory domain contains OUs, as outlined in the exhibit. The OU structure is designed around the company's remote office locations and all computer and server accounts (except for the local domain controllers and the local DHCP server which are found in the United States OU) are located in the OU named after where they are physically located.
- You need to ensure that the baseline security configuration for the systems in use at the CT OU complies with the written security policy that requires all traffic be encrypted between all systems at that location. You will be setting the security methods for main mode IKE negotiation for all of the systems on your network that require security. You need to specify which Diffie-Hellman groups to use on your systems. Which setting should be used to handle interoperability requirements between Windows 2000 and Windows XP systems?
- Select the best answer.
- A. Group 2 (medium) at 2048 bits
  - B. Group 2 (medium) at 1024 bits
  - C. Group 2048 (high) at 2048 bits
  - D. Group 2048 (high) at 4096 bits

**Exhibit(s):**

3. You are the domain administrator for gunderville.com. Client systems in use include Windows 2000 Professional and Windows XP Professional. Server systems in use include 2000 Server and Server 2003.

You have a number of group policies in use and they are linked at various objects in the Active-Directory hierarchy. You need a way to quickly assess the impact these linked policies and their configurations across the enterprise and well as review the details of group policy application. What new tool is included in Windows XP and Windows Server 2003 that allows you to verify the linked policies that have been applied and the outcome of the configuration settings on users and computers?

Select the best answer.

- A. Event Viewer
  - B. Network Monitor (NETMON)
  - C. SECDIT
  - D. Security Configuration and Analysis
  - E. Resultant Set of Policy (RSOP)
4. You are the enterprise administrator for gunderville.com. Client computers run Windows XP Professional and Windows 2000 Professional. You need to ensure that the baseline security configuration for the systems in use at the CT OU for your Active Directory structure complies with the written security policy, which requires that all traffic between systems be encrypted at that location.

You begin the process of editing your existing policy by double-clicking the policy and going to the General tab, then to Settings, and then choosing Methods.

In the Key Exchange Security Methods dialog box you need to modify an existing key exchange security method by clicking the security method that you want to modify, and then click Edit. Where would you need to go to configure the policy to use the SHA1 -160-bit key?

Select the best answer.

- A. Select from the available Integrity algorithms.
- B. Select from the available encryption algorithms.
- C. Select from the available Diffie-Hellman Groups.
- D. You would not be able to accomplish your task by editing an existing policy and editing the Key Exchange Security Methods.

## Chapter 6 Planning, Implementing, and Maintaining Security Infrastructure

1. You are the domain administrator for gunderville.com. Client systems in use include Windows 2000 Professional and Windows XP Professional. You have been asked to set up a Public Key Infrastructure configuration for your domain. You also need to identify which features will run in your environment.

Your domain controllers consist of only Windows Server 2003 systems and your Enterprise Certificate Authority is installed on a Windows Server 2003 Standard Edition system. Which of the following PKI features are supported in your environment?

Select all that apply.

- A. V2 templates
- B. Auto-enrollment for Computer certificates
- C. Auto-enrollment for User certificates
- D. Delta certificate revocation lists (CRLs)
- E. Role separation
- F. Qualified subordination

2. You are the domain administrator for your Windows Server 2003 mixed mode domain. There are 12 Windows Server 2003 systems in use across your enterprise, including 5 of the 12 installed as domain controllers.

Clients and servers in your domain consist of 73 Windows 98 systems, 6 Windows ME systems, 81 Windows NT4 Workstations running SP6a, 92 Windows 2000 Professional systems, 47 Windows 2000 Server systems, and 72 test Windows XP Professional systems with simple file sharing enabled.

You have been asked to perform an analysis of the systems in your environment by using the Microsoft Baseline Security Analyzer (MBSA). You realize that there are some client systems the tool will not be able to scan, but you are required to scan as many client systems as possible to report at the next security staff meeting.

You will be performing all of the scans remotely from your Windows Server 2003 Standard Edition system. Which client systems will you successfully be able to scan across the network?

Select the best answer.

- A. Windows NT4 Workstation
- B. Windows 98
- C. Windows ME
- D. Windows 2000 Professional
- E. Windows XP Professional



3. You are a domain administrator for gunderville.com with client systems running Windows 2000 Professional SP4 and Windows XP Professional SP1. You have been asked to configure the Windows Update clients in your DMZ to use a SUS version 1.0 server that has been installed in a segmented network. The name of the server is SUS001. Which types of client updates will be made available to the client and server operating systems via SUS001? Select all that apply.
- A. Designed for Windows Logo device drivers.
  - B. Windows 2000, Windows XP, and Windows Server 2003 Service Packs
  - C. Office 2000, Office XP, and Office 2003 Service Packs
  - D. Office 2000, Office XP, and Office 2003 Service Releases
  - E. Windows Update Rollups
  - F. Windows Security Patches (Critical, Important, Moderate, and Low)
4. You are the domain administrator for gunderville.com. As part of your CA design, your company has determined that it will be using Delta certificate revocation lists because of the benefit of the decreased network traffic that is caused when a new certificate revocation list needs to be downloaded and delta certificate revocation lists are used. What are two primary characteristics of delta certificate revocation lists? Select all that apply.
- A. Delta CRLs cannot be issued by Windows Server 2003 standalone Certificate Authorities.
  - B. Delta CRLs are issued by Windows Server 2003 standalone and enterprise CAs.
  - C. Delta CRLs are issued by Windows 2000 Server standalone and enterprise CAs.
  - D. Only clients that are running Windows XP Professional and later are able to check the validity of certificates against delta CRLs.

# Answers and Explanations

## Chapter 1

### 1. Answer: E

Explanation A. Although the Security Configuration and Analysis is a tool for analyzing and configuring local system security settings, it is a GUI based tool and would not be the easiest way to accomplish the required task because it would require you to run the tool on each system one at a time.

Explanation B. MBSA is a GUI tool (it can also be run from the command line) that allows an administrator to perform local or remote scans of Windows systems in an effort to scan for missing security updates and Service Packs for Windows, IE, IIS, SQL, Exchange, and Windows Media Player. It would not allow you to compare your standard desktop build's security settings against the default configuration security settings that are applied.

Explanation C. The File Signature Verification tool, SIGVERIF.exe, can be used to identify unsigned drivers on your system. It would not allow you to compare your standard desktop build's security settings against the default configuration security settings that are applied.

Explanation D. The System File Checker tool (SFC.exe) allows an administrator to scan all of the protected files on a computer to verify if they are the correct versions. It would not allow you to compare your standard desktop build's security settings against the default configuration security settings that are applied.

**Explanation E.** Secedit can be used to analyze and configure the security settings of computers by comparing your current configuration to at least one template from the command line. Using this tool as part of a script allows you to run it against all of the systems with less effort than most GUI tools.

### 2. Answers: C, E

Explanation A. The DnsAdmins group has administrative access to the DNS server service but the group has no default members and cannot be successfully used to complete your task. In order to start CMD.exe by using Secondary Logon (RUNAS), select Start, click Run, type runas /user:machine\_or\_DOMAIN \ADMIN\_ACCT cmd, (where machine\_or\_DOMAIN is the name of the local system or the domain, varying whether you are logging on locally or to the domain), and then click OK. A console window will appear, prompting for a password for the machine\_name\administrator account. Type the password for the administrator account and press Enter. A new console will appear running in the administrative context and any command-based programs can now be started from this console window and will be running in an administrative context.

Explanation B. Members of the Server Operators group can log on interactively to a DNS server and affect a certain level of administrative control over the system, but the group has no default members and cannot be successfully used to complete your task.

**Explanation C.** User accounts that have membership in the Domain Administrators group have the minimum level of required administrative privilege to perform the tasks required. Enterprise Administrators would also be able to perform these tasks but that would be a greater level of administrative control present than what was needed.

Explanation D. To enable secure dynamic updates for your DNS server, the DNS zone needs to be converted to Active Directory integrated, and as such, this is necessary administrative effort. This is not the case for standard dynamic updates, converting the zones is not necessary and is an additional administrative effort that is not necessary.

**Explanation E.** This is the proper syntax for you to ensure that dynamic updates are allowed for all of the specified DNS zones on the internal DNS servers. AllowUpdate 1 sets the configuration of the DNS zone specified in <ZONENAME> to allow dynamic updates. A setting of 0 would not allow dynamic updates and a setting of 2 configures the server to allow secure updates only, which would require that the zones were set up as Active Directory integrated.

Explanation F. This is not the proper syntax to ensure that dynamic updates are allowed for all of the specified DNS zones on the internal DNS servers. AllowUpdate 1 sets the configuration of the DNS zone specified in <ZONENAME> to allow dynamic updates. A setting of 0 would not allow dynamic updates and a setting of 2 configures the server to allow secure updates only, which requires that the zones were set up as Active Directory integrated.

3. Answers: C, E

Explanation A. On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

Explanation B. On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

**Explanation C.** On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

Explanation D. On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

**Explanation E.** On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

Explanation F. On a default installation of Windows XP Professional, all of the items listed under the Audit Policy of the default local security policy are configured with no auditing.

#### 4. Answer: A

**Explanation A.** To apply all of the security settings included in the security template, use the Security Configuration and Analysis MMC snap-in instead of the Local Computer Policy MMC snap-in. This is because the security settings for System Services cannot be applied using the Local Computer Policy MMC snap-in. Whenever you need to reset ACL permissions on a system back to the NTFS defaults or anytime you use the CONVERT command-line utility to change a FAT partition to NTFS (remember; the convert utility will set the ACLs for the converted drive to Everyone: Full Control), you need to run the security templates on the system in question as needed. For a Windows XP Professional workstation, this would be the defltwk.inf file; for a Windows Server 2003 member or stand alone system, this would be the defltsv.inf file; and for a Windows Server 2003 domain controller this, would be defltdc.inf.

Explanation B. To apply all of the security settings included in the security template you will need to use the Security Configuration and Analysis MMC instead of the Local Computer Policy MMC. This is because the security settings for System Services cannot be applied using the Local Computer Policy MMC snap-in.

Explanation C. With specific regard to default access control and local system services, whenever you need to reset ACL permissions on a system back to the NTFS defaults or anytime you use the CONVERT command-line utility to change a FAT partition to NTFS (remember; the convert utility will set the ACLs

for the converted drive to Everyone: Full Control), you need to run the security templates on the system in question as needed. For a Windows XP Professional workstation, this would be the defltwk.inf file; for a Windows Server 2003 member or stand alone system, this would be the defltsv.inf file; and for a Windows Server 2003 domain controller this, would be defltdc.inf.

Explanation D. The system in question is not a member of the domain. To apply all of the security settings included in the security template you will need to use the Security Configuration and Analysis MMC snap-in.

## Chapter 2

### 1. Answer: C

Explanation A. Each network segment can use either R4 or R5 because they have multiple paths to both. A link could go down or a single router could go down and any of the four subnets could still access the other three, the Internet, and all of the subnets at the company's main headquarters. The loss of the single link to the Windows Server 2003 Routing and Remote Access Server at IP address 172.15.40.254 or the loss of router R2 will cause the abrupt segmentation of the network and these two factors are the single points of failure of this network.

Explanation B. Adding these two routers on the 172.15.32 segment and the 172.15.47 segment would do nothing to alleviate the loss of the single link to the Windows Server 2003 Routing and Remote Access Server at IP address 172.15.40.254 or the loss of router R2, which would cause the abrupt segmentation of the network.

**Explanation C.** Because there is only one network segment from R2 to the Windows Server 2003 Routing and Remote Access Server at IP address 172.15.40.254, the loss of this segment would cause the abrupt segmentation of the network. The other single point of failure on this network is the R2 router because the loss of this router would also need to be addressed. Even if the additional network segment was added the failure of R2 would cause both of those segments to go offline.

Explanation D. This is the IP address of the Windows Server 2003 Routing and Remote Access Server and the question stated that you would need to alleviate the single points of failure on the network, if any, other than the hardware that makes up the Routing and Remote Access Server, which might isolate subnets from reaching other subnets of the Internet. The loss of the RRAS server hardware due to failure and the subsequent IP address bound to the NIC on that server should not be addressed as part of this assessment.

Explanation E. The loss of the single link to the Windows Server 2003 Routing and Remote Access Server at IP address 172.15.40.254 or the loss of router R2 would cause the abrupt segmentation of the network.

### 2. Answers: B, D

Explanation A. For the 149.88.72.0 range of IP addresses with a subnet mask of 255.255.255.0, the default gateway in use is fine.

**Explanation B.** This is one possible correct answer. The original IP address that was coming up during your PING test "SRV007.gunderville.com [149.88.72.19] with 32 bytes of data" was showing the IP address as 149.88.72.19 and IPCONFIG on the server itself showed an address of 149.88.72.91. Host 19 and 91 can exist on this subnet so this would allow SRV007 to connect to other systems. You would most likely be able to connect to SRV007 by IP address under this type of configuration. The DNS entry is most likely listing the original IP address of the old hardware configuration of SRV007 that was taken off the network and replaced with this system. The DNS attempts to the system are what are failing.

Explanation C. If this were the case, there would be trouble reaching other systems in this subnet and not just this one server. Additionally, the client systems would also be experiencing the issue.

**Explanation D.** This is one possible correct answer. The original IP address that was coming up during your PING test "SRV007.gunderville.com [149.88.72.19] with 32 bytes of data" was showing the IP address as 149.88.72.19 and IPCONFIG on the server itself showed an address of 149.88.72.91. Host 19 and 91 can exist on this subnet so this would allow SRV007 to connect to other systems. It may have been intentional to switch the host IP address but this information has not been updated in DNS according to what is shown.

Explanation E. This is not the case; 255.255.255 is a valid address to segmenting a class B address.

### 3. Answer: B

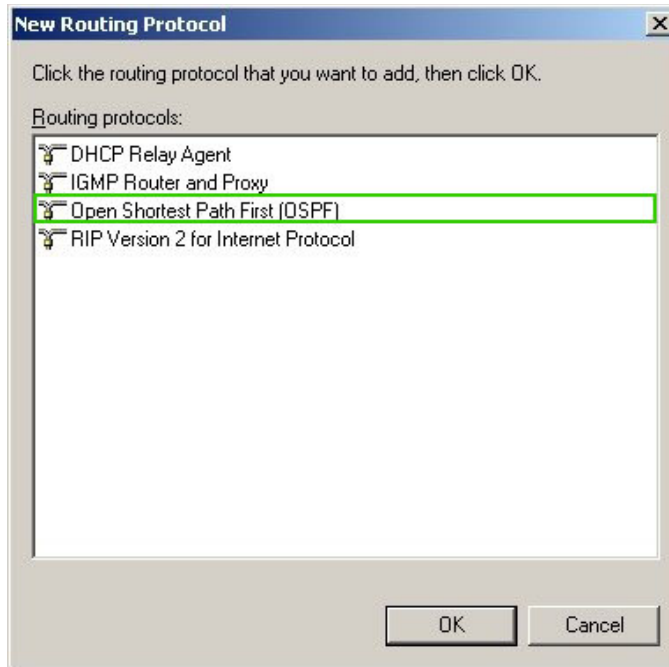
Explanation A. The IP address 18.12.23.253 is acceptable for the 18.12.16.0 subnetwork. When a subnet mask of 255.255.248.0 is used, any IP address from 18.12.16.1 to 18.12.23.254 would be acceptable as a host address. The main problem here is that the system was manually assigned a default gateway of 18.12.24.24, which is part of the next subnetwork - 18.12.24.1 through 18.12.31.254; with the remainder of the network information being supplied correctly (gateways, DNS servers, etc.) this causes communication issues with this system.

**Explanation B.** This is the best answer for this question. The main problem here is that the system was manually assigned a default gateway of 18.12.24.24, which is part of the next subnetwork - 18.12.24.1 through 18.12.31.254; with the remainder of the network information being supplied correctly (gateways, DNS servers, etc.) this causes communication issues with this system.

Explanation C. If this were the case, there would be trouble reaching other systems in this subnet and not just this one server.

Explanation D. The question stated that there are two WINS servers and two DNS servers local to each subnet and the clients are always configured to use those servers for name resolution only. The IP addresses for both the DNS servers and the WINS servers are correct for this subnet range.

#### 4. Answer:



**Explanation:** The OSPF protocol is the best choice when the network is designed with redundant paths between different locations or when the number of subnets in the overall design is more than 50 routers or when the two subnets are separated by 16 or more hops.

The 255.255.252.0 subnet mask is being used to allow for 1022 clients on a possible 64 subnets. Therefore, any dynamically updating update protocol needs to support Classless Inter-Domain Routing (CIDR) or variable-length subnet masks (VLSMs). RIPv1 does not allow for this. Also, RIP versions 1 and 2 are best used on medium-size networks with 50 routers maximum, and the maximum number of routers (hops) that any IP packet must cross is less than 16. Destination addresses that are 16 or more hops away are unreachable from RIP routers.

IGMP Router and Proxy is used by IP hosts to report their IP multicast group memberships to IGMP-enabled routers. It would not be used to update regular routing information.

**5. Answer:**

**Explanation:** You need to size the subnets for 450 clients per subnet (current) and also allow for an anticipated growth of 40% for the client systems, which is 630 clients. Therefore, you need to use the 255.255.252.0 subnet mask to allow for 1022 clients on a possible 64 subnets. Although this answer technically gives you too many subnet possibilities, it is the lowest number of hosts per subnet that works for the scenario.

A subnet mask length of 22 allows for 450 clients per subnet (current) and also allows for an anticipated growth of 40% for the client systems, which is 630 clients. You also need to use the 255.255.252.0 subnet mask to allow for 1022 clients on a possible 64 subnets.

A subnet mask length of 21 allows for 32 subnets but would yield 2046 hosts per subnet, which is more than needed. Although this configuration would work in the real world, it does not meet the requirements of the scenario.

A subnet mask of 255.255.248.0 allows for 32 subnets but would yield 2046 hosts per subnet, which is more than needed. Although this configuration would work in the real world, it does not meet the requirements of the scenario.

A subnet mask of 255.255.254.0 allows for 128 subnets, which is more than enough for the requirements in this scenario. This would yield 510 hosts per subnet. However, you need to size the subnets for 450 clients per subnet currently and also allow for an anticipated growth of 40% for the client systems (630). A subnet mask length of 23 allows for 128 subnets, which is more than enough for the requirements in this scenario. However, this would yield 510 hosts per subnet, which would not meet the expected 40% growth of client systems (to 630).

## Chapter 3

### 1. Answer: B

Explanation A. MS-CHAP v1 is a nonreversible, encrypted password authentication protocol that can be used where the authenticator sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string, but it cannot be used where a one-way encrypted password, mutual authentication process is used where the authenticator sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

**Explanation B.** MS-CHAP v2 provides a stronger security for remote access connections than MS-CHAP version 1 and it can be used where a one-way encrypted password, mutual authentication process is used where the authenticator sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

Explanation C. PEAP uses TLS to create an encrypted channel between an authenticating wireless computer and a PEAP authenticator. PEAP itself does not specify an authentication method, but provides additional security for other EAP authentication protocols that can operate through the TLS encrypted channel provided by PEAP.

Explanation D. EAP allows for open-ended conversation between the remote access client and the authenticating system. EAP is normally used with security token cards and smart cards where the authenticating system queries the remote access client for a name, PIN, and card token value or whatever else might be required during the challenge.

Explanation E. CHAP is a challenge-response authentication protocol that uses an MD5 hashing scheme to encrypt the response from the remote access client. It cannot be used where a one-way encrypted password, mutual authentication process is used where the authenticator sends a challenge to the remote access client that consists of a session identifier and an arbitrary challenge string.

### 2. Answer: B

Explanation A. The 255.255.255.128 subnet mask is being used to allow for 126 clients on a possible 512 subnets because 60 hosts per subnet are needed for now and the plans call for assuming growth of up to 69 to 70 hosts in total per subnet. This means that any dynamically updating update protocol should support CIDR or VLSM. RIPv1 does not allow for this.

**Explanation B.** This is the best answer. The 255.255.255.128 subnet mask is being used to allow for 126 clients on a possible 512 subnets because 60 hosts per subnet are needed for now and the plans call for assuming growth of up to 69 to 70 hosts in total per subnet. This means that any dynamically updating update protocol will need to support CIDR or VLSM. RIPv2 allows for this and this is the best answer for this question because Routing Information Protocol version 2 (RIP v2) is a distance vector routing protocol and provides both authentication and multicast announcing support.

Explanation C. Although the OSPF protocol is a better choice than either version of RIP when the network is designed with redundant paths between different locations, when the number of subnets in the overall design has more than 50 routers in use, or when remote locations are farther than 16 hops away. It is not a distance vector routing protocol; it is a link state routing protocol.

Explanation D. Interior Gateway Protocols (IGP) such as RIP or OSPF is used to exchange routing information within their networks. IGP is not something that could be used in the place of RIP or OSPF.



### 3. Answer: A

**Explanation A.** This is the correct answer. TTL boundaries prevent the forwarding of IP multicast traffic with a TTL less than a specified value and they apply to all multicast packets regardless of the multicast group. A TTL setting of 15 for a scope prevents the forwarding of IP multicast traffic that is intended to be restricted to the site. A higher time to live such as 63 for regional or 127 for worldwide for example would be routed forward.

Explanation B. This is not the correct answer as the TTL setting of 1 restricts the multicast traffic to the same subnet.

Explanation C. This is not the correct answer as the TTL setting of 63 restricts the multicast traffic to the same subnet region.

Explanation D. This is not the correct answer as the TTL setting of 127 forwards the multicast traffic worldwide.

Explanation E. This is not the correct answer as the TTL setting of 191 forwards the multicast traffic worldwide with limited bandwidth.

Explanation F. This is not the correct answer as the TTL setting of 255 forwards the multicast traffic totally unrestricted in scope.

### 4. Answer:

- A. Open the Configure your Server Wizard.
- B. Open the Routing and Remote Access MMC snap-in.
- C. Select General under Remote Access Policies and then right-click and choose New Routing Protocol.
- D. Select General under IP Routing, right-click, and choose New Routing Protocol.
- E. In the Select Routing Protocol dialog box, click Add, New Protocol, choose IGMP Router and Proxy, and then click OK.

B.  
D.  
E.

**Explanation:** To begin the process of configuring the Internet Group Management Protocol (IGMP), you would run the Routing and Remote Access MMC snap-in. Then, select General under IP Routing and click New Routing Protocol, which is part of configuring the Internet Group Management Protocol (IGMP) on your RRAS server. The last step of this process is to click IGMP Router and Proxy and then click OK to select IGMP for use.

Although the Configure your Server Wizard can be used to set up the Routing and Remote Access service initially to route on a LAN (or in other configurations such as a VPN server, etc.), it cannot be used to specifically allow you to configure the forwarding of IP multicast traffic and the use of the Internet Group Management Protocol (IGMP).

You cannot select General under Remote Access Policies.

## Chapter 4

### 1. Answers: A, B, E, F

**Explanation A.** Processor\% Processor Time is one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

**Explanation B.** Network Interface Connection\Bytes Total/sec is one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

**Explanation C.** This is not one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

**Explanation D.** This is not one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

**Explanation E.** PhysicalDisk\% Disk Time is one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

**Explanation F.** PhysicalDisk\Current Disk Queue Length is one of the four counters that should be collected in an effort to determine the frequency of disk access. If all four of the recommended counters have high values, the hard disk is not causing a bottleneck. If the % Disk Time is high and the processor and network connection are not saturated, the hard disk may be creating a bottleneck. The Current Disk Queue Length value should not be greater than 2.

### 2. Answer: D

**Explanation A.** Microsoft Message Queuing (MSMQ) is used to provide guaranteed message delivery, efficient routing, security, and priority-based messaging. It cannot be used to extend Cluster Shared Disk Partitions on your configuration as needed without disrupting the current data on the volumes.

**Explanation B.** Although it might be possible for you to extend the size for the current volume by using the Disk Management tool within the Computer Management MMC, it would not be the best way to do it and not disrupt the current data on the volumes.

**Explanation C.** Although it might be possible for you to extend the size for the current volume by using the Disk Management MMC, it would not be the best way to do it and not disrupt the current data on the volumes.

**Explanation D.** This is the best answer to this question because DISKPART.EXE allows you to extend the size for the current volume as long as the underlying storage hardware supports dynamic expansions of a disk unit (LUN -Logical Unit Number).

3. Answer: D

Explanation A. Nothing additional needs to be done except replacing the failed drives because the RAID configurations of this system will allow the system to continue to run intact.

Explanation B. Nothing additional needs to be done except replacing the failed drives because the RAID configurations of this system allow the system to continue to run intact.

Explanation C. Nothing additional needs to be done except replacing the failed drives because the RAID configurations of this system allow the system to continue to run intact.

**Explanation D.** When disk1 fails, it is one disk of a mirror (RAID1) set; the partition information and all of the data on it is still available to the system on the other drive (disk0) in that set. Once this failure occurs, disk2 comes online and will be built up with the data information via the server's hardware controller and the parity information found on disk0 which is still running. Once this is completed, the RAID1 mirror and the redundancy will be restored. When the drive in the RAID 5 configuration fails (disk3), the system is able to continue to function normally from the stored parity information so all of the data in the RAID5 array is still accessible. When disk2 fails, all of the data on the first array is still available because the fully rebuilt RAID1 mirror can suffer the single loss of a drive again. All of the redundancy on this system is lost but it is still 100% accessible.

#### 4. Answer: A

**Explanation A.** Given that the currently installed version of the server OS is Windows Server 2003 Enterprise Edition, you could scale out this cluster configuration to eight servers with eight CPUs and 32GB of RAM installed in each system at a maximum.

Explanation B. Given that the currently installed version of the Server Operating system is Server 2003 Enterprise Edition, you could scale out this cluster configuration to eight servers with eight CPUs and 32GB of RAM installed in each system at a maximum.

Explanation C. Given that the currently installed version of the Server Operating system is Server 2003 Enterprise Edition, you could scale out this cluster configuration to eight servers with eight CPUs and 32GB of RAM installed in each system at a maximum.

Explanation D. Given that the currently installed version of the Server Operating system is Server 2003 Enterprise Edition, you could scale out this cluster configuration to eight servers with eight CPUs and 32GB of RAM installed in each system at a maximum.

## Chapter 5

### 1. Answers: C, F

Explanation A. The Windows 98 clients are in branch office one which is connected to the company headquarters by a private leased line; it does not connect via the Internet, so there is no requirement to install the Microsoft L2TP/IPSec VPN Client.

Explanation B. The RRAS server itself would be making the connection and the security association; the clients themselves would pass through the tunnel created by the RRAS server and the only issue is at branch office one which is actually negated by the fact that connection is made by a private leased line.

**Explanation C.** The Windows 98 clients are in Branch office one which is connected to the company headquarters by a private leased line; it does not connect via the Internet so there is no requirement to install the Microsoft L2TP/IPSec VPN Client.

Explanation D. The Microsoft L2TP/IPSec VPN Client would need to be installed on all Windows 98 systems making the connections to the VPN server individually by way of the Microsoft L2TP/IPSec VPN adapter, but this is not needed in this scenario because the connection from branch office one is via a private leased line. If it was necessary, the PPTP adapter would be available for MPPE (Microsoft Point-to-Point Encryption) secured transmissions.

Explanation E. ESP is more than what is required here. You need to ensure that all of this data is secured against tampering and unauthorized modification. Although ESP does do this, it also encrypts the data. This adds a layer of security but it is an unneeded layer that adds additional resource load to the servers or networks than is absolutely necessary.

**Explanation F.** You needed only to ensure that all of this data is secured against tampering and unauthorized modification; AH does exactly this.

### 2. Answer: B

Explanation A. Group 2 (medium) operates at 1024 bits not 2048 as the answer suggests.

**Explanation B.** This is the correct answer. Diffie-Hellman groups are used to determine the length of the base prime numbers used during the key exchange process. The cryptographic strength of any key is partially tied to the strength of the Diffie-Hellman group upon which the prime numbers are based. Group 2048 (high) is stronger than Group 2 (medium) at 1024 bits, which is stronger than Group 1 (low) at 768 bits of keying strength. If there are mismatched Diffie-Hellman groups specified on different systems that require security, the secure association negotiation will fail. With regard to interoperability requirements between Windows 2000 and Windows XP systems, you should opt to use Group 2 (medium) at 1024 bits. Currently, Diffie-Hellman Group 2048 is provided only with Windows Server 2003 operating systems.

Explanation C. This is not the correct answer. With regard to interoperability requirements between Windows 2000 and Windows XP systems, you should opt to use Group 2 (medium) at 1024 bits. Currently, Diffie-Hellman Group 2048 is provided only with Windows Server 2003 operating systems.

Explanation D. Group 2048 (high) runs at 2048 bits and not 4096 as the proposed answer suggests. With regard to interoperability requirements between Windows 2000 and Windows XP systems, you should opt to use Group 2 (medium) at 1024 bits. Currently, Diffie-Hellman Group 2048 is provided only with Windows Server 2003 operating systems.

### 3. Answer: E

Explanation A. Although the Event Viewer can be used to confirm the application of group policy, it cannot give you the detailed output that is required here. Also, the Event Viewer is not a new tool.

Explanation B. Network Monitor can be used to view network traffic. It cannot be used to give you the detailed output on group policy objects that is required here. Also, the Network Monitor is not a new tool.

Explanation C. SECEDIT can be used to configure and analyze system security by comparing your current configuration to at least one template but it is not a new tool to the Windows XP and Windows Server 2003 operating systems.

Explanation D. You can use Security Configuration and Analysis to analyze and configure local computer security but it is not a new tool to the Windows XP and Windows Server 2003 operating systems.

**Explanation E.** The Resultant Set of Policy (RSOP) is a new tool to the Windows XP and Windows Server 2003 operating systems that allows you to verify policies in effect for different users or computer so that you can quickly assess configurations throughout your enterprise. RSOP can be configured so that the MMC snap-in is directed to check policies for any computer or user in a domain.

### 4. Answer: A

**Explanation A.** From the IKE Security Algorithms dialog box, select one of the available Integrity algorithms which are: MD5 (uses a 128-bit key) or SHA1 (uses a 160-bit key).

Explanation B. You can choose either 3DES to use the triple Data Encryption Standard (3DES) algorithm and three unique 56-bit keys or DES to use the DES algorithm and a single 56-bit key. You cannot select the SHA1 Integrity algorithm from here.

Explanation C. You would select a Diffie-Hellman group to set the length of base keying material used to generate the actual keys; you cannot select the SHA1 Integrity algorithm from here.

Explanation D. This is not the correct answer as these beginning steps as outlined in the question are the correct ones to perform.

## Chapter 6

### 1. Answers: B, D, F

Explanation A. V2 templates are supported only under Windows Server 2003 Enterprise Edition or Datacenter Edition when configured as an Enterprise CA.

**Explanation B.** Auto-enrollment for Computer certificates is supported under Windows Server 2003 Standard Edition Certificate Authorities.

Explanation C. Auto-enrollment for Computer certificates is not supported under Windows Server 2003 Standard Edition Certificate Authorities. Both user and computer certificates are supported under Windows Server 2003 Enterprise Edition or Datacenter Edition.

**Explanation D.** Delta certificate revocation lists (CRLs) are supported under Windows Server 2003 Standard Edition Certificate Authorities as well as under Windows Server 2003 Enterprise Edition or Datacenter Edition.

Explanation E. Role separation is not supported under Windows Server 2003 Standard Edition Certificate Authorities. Role separation is supported under Windows Server 2003 Enterprise Edition or Datacenter Edition.

**Explanation F.** Qualified subordination is supported under Windows Server 2003 Standard Edition Certificate Authorities as well as under Windows Server 2003 Enterprise Edition or Datacenter Edition.

## 2. Answers: A, D

**Explanation A.** MBSA Version 1.1.1 can perform a local or network scan of systems running Windows Server 2003, Windows 2000, and Windows XP Professional. (XP Pro is affected by simple file sharing if it is enabled on a system; then only a local scan can be performed. A network scan will be prohibited.) In addition, MBSA can scan Windows NT 4.0 SP4 and above systems over the network. Windows 98 and/or ME systems cannot be scanned by the tool.

Explanation B. MBSA Version 1.1.1 can perform a local or network scan of systems running Windows Server 2003, Windows 2000, and Windows XP Professional. (XP Pro is affected by simple file sharing if it is enabled on a system; then only a local scan can be performed. A network scan will be prohibited.) In addition, MBSA can scan Windows NT 4.0 SP4 and above systems over the network. Windows 98 and/or ME systems cannot be scanned by the tool.

Explanation C. MBSA Version 1.1.1 can perform a local or network scan of systems running Windows Server 2003, Windows 2000, and Windows XP Professional. (XP Pro is affected by simple file sharing if it is enabled on a system; then only a local scan can be performed. A network scan will be prohibited.) In addition, MBSA can scan Windows NT 4.0 SP4 and above systems over the network. Windows 98 and/or ME systems cannot be scanned by the tool.

**Explanation D.** MBSA Version 1.1.1 can perform a local or network scan of systems running Windows Server 2003, Windows 2000, and Windows XP Professional. (XP Pro is affected by simple file sharing if it is enabled on a system; then only a local scan can be performed. A network scan will be prohibited.) In addition, MBSA can scan Windows NT 4.0 SP4 and above systems over the network. Windows 98 and/or ME systems cannot be scanned by the tool.

Explanation E. MBSA Version 1.1.1 can perform a local or network scan of systems running Windows Server 2003, Windows 2000, and Windows XP Professional. (XP Pro is affected by simple file sharing if it is enabled on a system; then only a local scan can be performed. A network scan will be prohibited.) In addition, MBSA can scan Windows NT 4.0 SP4 and above systems over the network. Windows 98 and/or ME systems cannot be scanned by the tool. The question stated that the Windows XP Professional computers were test systems with simple file sharing enabled. Because of this setting, you would not be able to scan them over the network.

## 3. Answers: B, E, F

Explanation A. Although these updates are available on the Windows Update site, you cannot use SUS to distribute them.

**Explanation B.** The following types of updates are currently supported for SUS 1.0:

- Windows Critical Updates
- Windows Security Patches (Critical, Important, Moderate, and Low)
- Windows Update Rollups

- Windows 2000, Windows XP, and Windows Server 2003 Service Packs

Explanation C. SUS 1.0 cannot be used to distribute Office Service Packs and neither can the Windows Update site at this time. The future configurations of Windows Update (to be called Microsoft Update) and SUS are being planned to provide this functionality.

Explanation D. SUS 1.0 cannot be used to distribute Office Service Releases and neither can the Windows Update site at this time. The future configurations of Windows Update (to be called Microsoft Update) and SUS are being planned to provide this functionality.

**Explanation E.** The following types of updates are currently supported for SUS 1.0:

- Windows Critical Updates
- Windows Security Patches (Critical, Important, Moderate, and Low)
- Windows Update Rollups
- Windows 2000, Windows XP, and Windows Server 2003 Service Packs

**Explanation F.** The following types of updates are currently supported for SUS 1.0:

- Windows Critical Updates
- Windows Security Patches (Critical, Important, Moderate, and Low)
- Windows Update Rollups
- Windows 2000, Windows XP, and Windows Server 2003 Service Packs

#### 4. Answers: B, D

Explanation A. Delta CRLs are available for use under Windows Server 2003 standalone and enterprise CAs.

**Explanation B.** This is a correct answer; Delta CRLs are issued by Windows Server 2003 standalone and enterprise CAs.

Explanation C. Windows 2000 Server CAs cannot use delta CRLs. Also only clients that are running Windows XP Professional and later are able to check the validity of certificates against delta CRLs on Windows Server 2003 standalone and enterprise CAs.

**Explanation D.** This is a correct answer; clients that are running Windows XP Professional and later are able to check the validity of certificates against delta CRLs.