

Microsoft

Server 2003

Designing Network Security

(70-298)

Microsoft Certified
Systems Engineer (MCSE)



**Smarter
Training**

This LearnSmart exam manual covers the most important concepts you need to know in order to successfully complete the Server 2003 Designing Network Security exam (70-298). By studying this exam manual, you will become familiar with a variety of exam-related content, including:

- Analyzing Business Requirements
- Analyzing Technical Requirements
- Designing a Windows 2000 Security Solution
- And more!

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

Server 2003 Designing Network Security (70-298) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC
Product ID: 10019
Production Date: July 12, 2011
Total Questions: 26

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

| | |
|---|-----------|
| Abstract | 6 |
| What to Know | 6 |
| Tips | 7 |
| Creating the Conceptual Design for Network Infrastructure Security | 8 |
| Analyzing Business Requirements for Designing Security | 8 |
| <i>Business Factors that Dictate Security Designs</i> | 8 |
| <i>Analyzing Existing Security Policies and Procedures</i> | 9 |
| <i>Analyzing the Requirements for Data Security</i> | 9 |
| <i>Analyzing Security Risks within the IT Administration Structure</i> | 11 |
| Designing a Security Implementation Framework | 11 |
| <i>Predicting Internal and External Threats to Your Network</i> | 11 |
| <i>Designing an Incident Response Process</i> | 12 |
| <i>Designing Segmented Networks</i> | 13 |
| <i>Designing a Service Recovery Process</i> | 13 |
| Technical Constraints in Security Design | 14 |
| <i>Identifying Capabilities of the Existing Infrastructure</i> | 14 |
| <i>Identifying Technology Limitations</i> | 15 |
| <i>Analyzing Interoperability Constraints</i> | 15 |
| Creating the Logical Design for Network Infrastructure Security | 16 |
| Designing a Public Key Infrastructure using Certificate Services | 16 |
| <i>Designing a Certification Hierarchy Implementation</i> | 16 |
| <i>Designing Enrollment and Distribution Processes</i> | 18 |
| <i>Certificate Renewal</i> | 19 |
| <i>Certificate Revocation</i> | 19 |
| <i>Certification Authority Auditing</i> | 20 |
| <i>Designing Security for Certificate Servers</i> | 21 |
| Designing a Logical Authentication Strategy | 21 |
| <i>Designing Certificate Distribution</i> | 21 |
| <i>Designing Forest and Domain Trust Models</i> | 22 |
| <i>Designing Security that Meets Interoperability Requirements</i> | 23 |
| <i>Establishing Account and Password Requirements for Security</i> | 23 |
| Designing Security for Network Management | 24 |
| <i>Managing the Risk of Managing Networks</i> | 24 |

| | |
|--|-----------|
| <i>Designing Server Administration by Using Common Administrative Tools</i> | 25 |
| <i>Designing Security for Emergency Management Services</i> | 26 |
| Designing a Security Update Infrastructure | 27 |
| <i>Designing a Software Update Services (SUS) Infrastructure</i> | 27 |
| <i>Using Group Policy to Enable Deployment of Software Updates</i> | 28 |
| <i>Identifying Computers that are not at the Current Patch Level</i> | 29 |
| Creating the Physical Design for Network Infrastructure Security | 30 |
| Designing Network Infrastructure Security | 30 |
| <i>Specifying Firewall Protocols</i> | 30 |
| <i>Specifying IP Packet Filtering</i> | 31 |
| <i>Designing Internet Protocol Security (IPSec) Policies</i> | 32 |
| <i>Securing a Domain Name Service (DNS) Implementation</i> | 35 |
| <i>Designing Security for Data Transmission</i> | 36 |
| Designing Security for Internet Information Services | 36 |
| <i>Designing IIS User Authentication</i> | 37 |
| <i>Use of RADIUS for IIS Authentication</i> | 38 |
| <i>Use of Certificates for IIS Authentication</i> | 38 |
| <i>Designing IIS Security Baselines According to Business Needs</i> | 38 |
| <i>Designing Web Site Security By Enabling the Minimum Required Services</i> | 39 |
| <i>Designing a Monitoring Strategy for IIS</i> | 41 |
| <i>Designing a Content Management Strategy for Updating IIS Servers</i> | 42 |
| Designing Security for Communication Between Networks | 42 |
| <i>Selecting Protocols for VPN Access</i> | 42 |
| <i>Designing VPN Connectivity</i> | 43 |
| <i>Designing Demand-Dial Routing Between Internal Networks</i> | 44 |
| Designing Security for Wireless Networks | 45 |
| <i>Designing Public and Private Wireless LANs</i> | 45 |
| <i>Designing 802.1x Authentication for Wireless Networks</i> | 46 |
| Designing Security for Communication with External Organizations | 48 |
| <i>Designing an Extranet Infrastructure</i> | 48 |
| <i>Designing a Strategy for Cross-Certification of Certificate Services</i> | 48 |
| Designing Security for Servers with Specific Roles | 49 |
| <i>Defining a Baseline Security Template for All Systems</i> | 49 |
| <i>Creating a Plan to Modify Baseline Security Templates According to Role</i> | 51 |

| | |
|---|-----------|
| Designing an Access Control Strategy for Data | 52 |
| Designing an Access Control Strategy for Files and Folders | 52 |
| <i>Designing a Permission Structure for Files and Folders</i> | 52 |
| <i>Designing an Encryption and Decryption Strategy</i> | 53 |
| <i>Designing Security for Backup and Recovery</i> | 54 |
| Designing an Access Control Strategy for Directory Services | 55 |
| <i>Creating a Delegation Strategy</i> | 55 |
| <i>Designing a Group Strategy for Accessing Resources</i> | 57 |
| <i>Designing a Permission Structure for Directory Service Objects</i> | 57 |
| Designing an Access Control Strategy for the Registry | 58 |
| Analyzing Auditing Requirements for | |
| Directory Services, Files and Folders, and the Registry | 60 |
| <i>Designing Auditing for Directory Service Objects</i> | 61 |
| <i>Designing Auditing for Files and Folders</i> | 63 |
| <i>Designing Auditing for the Registry</i> | 63 |
| Creating the Physical Design for Client Infrastructure Security | 64 |
| Designing a Client Authentication Strategy | 64 |
| <i>Analyzing Authentication Requirements</i> | 64 |
| <i>Establishing Account and Password Security Requirements</i> | 65 |
| Designing a Security Strategy for Client Remote Access | 65 |
| <i>Designing Remote Access Policies</i> | 65 |
| <i>Designing Access to Internal Resources</i> | 67 |
| <i>Using IAS to Provide Authentication and Accounting for Remote Network Access</i> | 69 |
| Designing a Strategy for Securing Client Computers | 70 |
| <i>Designing a Strategy for Hardening Client Operating Systems</i> | 70 |
| <i>Designing a Strategy for Restricting User Access to Operating</i> | 74 |
| Practice Questions | 76 |
| Answers and Explanations | 91 |

Abstract

This Exam Manual will help you prepare for the Microsoft 70-298, Designing Security for a Microsoft Windows Server 2003 Network exam. Exam topics include five domains. The exam assesses your ability to design security solutions in a Windows Server 2003 environment. Topics covered include gathering and analyzing data in support of business and technical requirements for network security; creating the logical design for network security, including public key infrastructure (PKI), authentication, network management, and security updates; creating the physical design for network security, including wireless networks, Internet Information Services (IIS), communication between networks and with external organizations, and security for servers with specific roles; designing an access control strategy for data, including directory services, files, folders, and the registry; creating the physical design for client infrastructure security, including client authentication and remote access; and securing desktop and portable client computers.

What to Know

Be sure to familiarize yourself with the following concepts:

- Analyzing business requirements for security design
- Realizing technical constraints and limitations for security design
- Designing a PKI that uses Certificate Services
- Designing a logical authentication strategy
- Designing network management security
- Designing a security update infrastructure
- Designing network infrastructure security including wireless network security
- Designing user authentication and security for IIS
- Designing security for communication between networks and with external companies
- Designing security for servers with specific roles
- Designing access control strategies for directory services, files and folders, and the registry
- Designing strategies for client authentication and client remote access
- Designing strategies for securing desktop and portable client computers

Make sure to take time to review the exam objectives at [Microsoft](#).

Tips

This exam is case-study-based. It consists of a series of testlets in which you are provided information on a security scenario and required to answer 8 to 12 questions based on the facts presented. You will have 2 hours to complete it. It covers questions on designing security on Windows Server 2003 computer networks.

In preparing for the exam, you should work with the tools and techniques covered on the exam. In particular, you should set up a network of computers running Windows Server 2003 and Windows XP Professional, and install and configure tools such as Certificate Services, Terminal Server, Remote Desktop, Remote Assistance, Software Update Services (SUS), Microsoft Baseline Security Analyzer (MBSA), Internet Information Services (IIS), Routing and Remote Access (RRAS), and Internet Authentication Service (IAS). You should practice with the techniques tested on the exam, including client authentication, access control, IIS authentication and access control, Internet Protocol Security (IPSec), including IP filtering, Domain Name Service (DNS), security updates and patches, certificate enrollment and distribution, and so on.

You can obtain a 180-day evaluation version of Windows Server 2003 from Microsoft. You can also use virtual computer software, such as Microsoft Virtual PC, to set up multiple virtual computers on a single physical machine. Evaluation versions of virtual computer software are also available.

With two hours reserved for the exam, you will have plenty of time to complete the test. Don't rush it. Review the facts presented in each testlet carefully and make notes of the most important facts on the plastic sheets provided. You can toggle back and forth between the case study facts and the questions as much as needed. However, once you've completed a testlet, you cannot return to it later. On the first pass through a testlet, work through the questions and answer all that you are sure of and then on the second pass, spend time on the questions that presented more of a challenge. Even if you are unsure of the correct answer, determine the ones you know are incorrect. This helps you narrow your options and gives you a better chance at passing the exam.

Creating the Conceptual Design for Network Infrastructure Security

News events relating compromise of important data have been around for years, but people have become more aware of data security and potential for problems such as identity theft in recent years, and especially since September 11, 2001. But security risks come from sources large and small, even from employees within your organization who you thought were highly trustworthy. For an overview of designing a secure Windows Server 2003 environment, consult links contained in the [Windows Server Deployment Manual](#). Another excellent source of information related to all objectives of the 70-298 exam is Microsoft's [Windows Server 2003 Security Manual](#). This is a valuable companion to this manual for anyone desiring to pass the 70-298 exam.

Analyzing Business Requirements for Designing Security

Business managers look to the IT department as a service that keeps their business operating at proper efficiency and produces the highest level of profit for their company. They do not always realize that information security is part and parcel of a proper business design. The very existence of a company could be threatened should information, such as financial data, research and development, legal, and so on, fall into the wrong hands. Refer to [Identifying Business Design Requirements](#) for an introduction to several design requirements that you should consider.

Business Factors that Dictate Security Designs

As a network security designer, you must collaborate with executives and managers who are responsible for setting business requirements. It is essential that you communicate the need for proper network security and the potential consequences of inadequate protection of network assets. Refer to [Designing for Securability](#) for an overview of several threats to network security and their associated risks. Factors to address include the following:

- *Policies and procedures* – Written security policies stipulate the steps that must be followed to ensure the security of company networks and data, as well as the consequences to employees of violating these policies.
- *Sensitivity of data* – Not all data needs the same level of security. Whereas some data may be designed for public disclosure, other types of data may be restricted for employees only, and still other data is confidential to managers only. See [Business Rules for Data Access](#) for an overview of several factors affecting data security.
- *Cost* – Should a security breach occur, you will incur costs associated with determining the source of the breach, the amount of damage that has occurred, and so on. Costs will be associated with restoring systems and data, and even with restoring consumer confidence should your Web site be compromised and customer information end up in the wrong hands. Balance all such costs against the cost of designing and implementing an adequate security solution in your organization.
- *Legal needs* – New laws in the United States and elsewhere dictate requirements that impact security design. You should enlist the cooperation of company legal staff and (where necessary) outside legal help when completing your network security design. The following are several U.S. laws that influence security design:
 - The Health Insurance Portability and Accountability Act (HIPAA) governs security requirements for organizations involved in health care. See [Draft HIPAA Security Summit Guidelines](#) for security guidelines.

- ▶ The Graham-Leach-Bliley Act mandates the security and confidentiality of customer financial data. See [What is the Graham-Leach Bliley Act of 1999?](#) for an overview of this law.
- ▶ The Sarbanes-Oxley Act of 2002 mandates requirements for business and financial data. See [What the Sarbanes-Oxley Act means for IT managers](#) for an overview of this law's implications on IT management and design.
- *End-user impact* – The higher the level of security you implement on the network, the greater the impact on the end users. For example, a strong password policy increases the chance that users will write down passwords on sticky notes and attach them to their monitor or keyboard. Strong restrictions on Internet access increase the chance that users will install unauthorized modems. Account lockout policies increase the number of help desk calls and reduce productivity when users are unable to access computers and networks. See [Privacy and Security](#) for more information.
- *Interoperability, Maintainability, and Scalability* – All networks contain various components that must work together to provide reliability and security. Inventory the types of devices, such as firewalls, routers, and servers, as well as the software and network protocols, in use. Pay particular attention to older systems, such as Windows NT and 9x, that do not support the latest security technologies.
- *Risk* – A risk can be defined as the probability of an undesirable effect resulting from a threat being realized. Although it is possible to prepare a quantitative risk assessment that places dollar amounts on each risk, an evaluation of the qualitative impact of risks is often sufficient. See [Planning for Deployment](#) for more information.

Analyzing Existing Security Policies and Procedures

Security policies define the types of activities that are allowed or forbidden on a company's network and with company-owned devices, such as desktops, laptops, personal data assistants (PDAs), and so on. Companies should have a written security policy that has been reviewed and approved by management and the company legal staff. *Security procedures* refer to the types of actions undertaken to achieve compliance with security policies.

Section 2 of the [Site Security Handbook](#) describes typical components of written security policies. Also refer to [Security Strategies](#) for examples of reviewing current security policies and methodologies for reviewing policies and defining strategies. Keep in mind that case studies on the 70-298 exam contain written security policies that you must analyze to obtain information needed for answering questions correctly.

Analyzing the Requirements for Data Security

Companies must ensure that their data is secured so that its confidentiality, integrity, and availability are guaranteed. Different types of organizations have different requirements for data security, and different types of data need to be secured in different manners. [Service Management Functions](#) reviews needs for data security and types of defenses you can put into place for protecting your data.

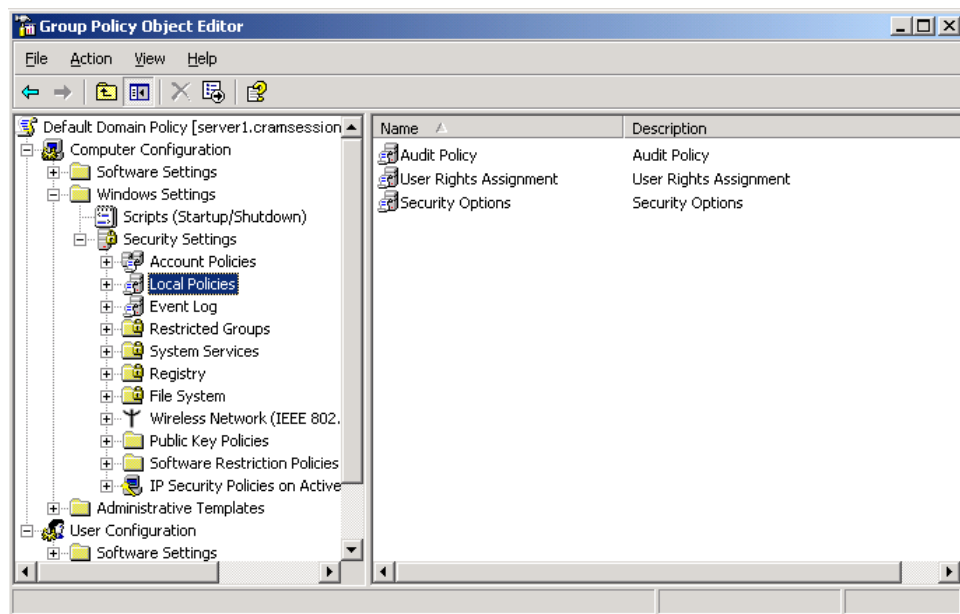
You should take the following considerations into account when designing a data security strategy:

- *The purpose and sensitivity of the data* – Certain types of data, such as accounts receivable, need a higher security level than other data types, such as an employee telephone list.
- *The impact of errors in the data* – You must ensure data integrity. Think of what would happen if customers received bills for a small percentage of the amount actually owed; how long could your company survive?

- *The trust placed in your organization by outside parties* – Your customers must be assured that their confidential information, such as credit card numbers, will not be compromised or divulged to unauthorized parties.

Windows Server 2003 and key Microsoft applications contain numerous components that you can utilize for securing your data. The following list is not intended to be exhaustive:

- [Group Policy](#) (Figure 1) enables you to define security settings to be applied to a computer, user, or groups of computers or users.



- [Security Configuration Manager](#) includes several tools that you can use to configure security settings for computers, organizational units (OUs), and domains. It includes [Security Templates](#), which enables you to define security settings that can be applied to computers using group policy or the [Security Configuration and Analysis](#) tool.
- [Auditing](#) enables you to audit access to data and other components so that you can track usage and spot unauthorized attempts to access or modify data.
- Microsoft Office 2003 contains security features that enable you to better secure your data. For additional information, consult [Microsoft Office 2003 Editions Security Whitepaper](#).

Analyzing Security Risks within the IT Administration Structure

Designing an adequate security infrastructure for your organization requires an understanding of the inherent risks to the security of your data. Evaluate the administrative structure of your organization and the IT department so that managers and other important individuals are aware of the risks and the need to take the appropriate remedial measures. The [Microsoft Operations Framework Risk Management Discipline](#) presents the principles of a risk management and evaluation strategy that you can use to perform a risk analysis.

Designing a Security Implementation Framework

A security implementation framework is a structure upon which you can build components that assist you in addressing threats and vulnerabilities to your network and the means of dealing with them. Concepts you should include when designing the framework include the following:

- *Prevention* – By following a proactive approach to security design, you can prevent many problems from happening. A good example is the use of antivirus, antispam, and other similar software at the enterprise level. The use of [ISA Server 2004](#) to secure the network perimeter is another good example.
- *Detection* – You must detect any attacks in progress as early as possible to minimize damage. Intrusion detection systems (IDSs) are designed to perform this type of job. These are designed to operate on the entire network or specific hosts within the network, and can send messages to the administrator or even shut down the network or segments thereof. For questions and answers about IDSs refer to [Intrusion Detection FAQ](#).
- *Isolation* – You should isolate computers affected by an attack to prevent its progress and gather information that will lead to the identification of the attacker and prevent further intrusions. Many types of IDSs can perform this action.
- *Recovery* – By following an adequate backup and restore procedure that includes all types of data and software in use on your network, you can ensure that recovery from an attack can be accomplished rapidly with minimal loss of productivity.

Predicting Internal and External Threats to Your Network

To understand and predict threats to your network, you must have an idea of the vulnerabilities to which the network is subject. Review [Computer Vulnerabilities](#) for an overview of the existing types of threats. Threats to your network can come from the following sources:

- *Internal* – Activities taking place within your own network are often the most severe risks to the network. These can range from simple curiosity ("What can I find on My Network Places?") to vandalism perpetrated by a disgruntled or terminated employee.
- *External* – Attacks originating from outside the network that pass through an inadequate or nonexistent firewall. According to information provided by the [SANS Institute](#), an unsecured computer connected to the Internet will be compromised within 15 to 20 minutes.

The process of *threat modeling* serves to identify and predict threats and vulnerabilities to computers and networks. It involves the following [components](#):

1. *Identifying the assets* – Determining the most valuable network components that need the highest level of protection.
2. *Creating an architecture* – Documenting the structure of your networks and applications.
3. *Decomposing the application* – Creation of security profiles of applications running on your network that identify vulnerabilities within their design, deployment, or implementation.
4. *Identifying the threats* – Thinking in terms of what an attacker might be seeking, decide what types of threats you must protect your network and applications against. Microsoft recommends the use of the [STRIDE](#) model as a framework for threat identification.
5. *Documenting the threats* – Provide details about each threat you have identified and how it can impact upon your network and its daily operations.
6. *Rating the threats* – Prioritization of the threats and the potential extent of damage each may inflict. Doing so enables you to direct your preventative actions toward the most important threats.

Designing an Incident Response Process

No matter how comprehensive your risk analysis and security prevention plans are, some type of incident is bound to occur sooner or later. When an incident does occur, it is essential that you have a well thought-out response procedure in place. A successful [incident response procedure](#) should include the following components:

- *Efforts to minimize the number and severity of incidents* – These include establishment and enforcement of policies and procedures, obtaining management support, assessment of vulnerabilities, security training for IT staff and end users, information banners, implementation of password policies, network and system performance monitoring, analysis of logs, and verified backup procedures.
- *Use of a Computer Security Incident Response Team (CSIRT)* – The CSIRT should include individuals with responsibility for dealing with all types of security incidents. These individuals need the appropriate training and tools (including communication capabilities) so that they can respond rapidly to the first evidence of an intrusion.
- *Definition of an incident response plan* – All users should be aware of symptoms that may indicate that an intrusion is taking place. Furthermore, they need to be aware of how to communicate their observations to IT staff, managers, and the CSIRT. The response plan needs components such as communications procedures, containment of the incident, notification of external agencies as appropriate, documentation of the incident, assessment of damage and mitigation procedures, and review procedures in place.
- *Containment of the damage and minimization of risk* – Rapid action is necessary to protect your systems and network and prevent a minor incident from escalating into a major one. You should place priority on certain actions, such as protection of users' safety, data protection (especially sensitive data, but do not neglect other data including scientific, management, legal, and proprietary data), protection of hardware and software against further attack, and the minimization of disruption to computing resources.

Designing Segmented Networks

A segmented network refers to a network containing one or more devices designed to limit the data flow between network segments, such as firewalls, routers, and so on. Check out [Security Zones](#) for details on the following three types of segmented networks:

- *Bastion host* – This is a server or router with at least two network interface cards (NICs), which acts as a firewall to separate an internal network segment from a potentially hostile external network, such as the Internet.
- *Screened host gateway* – This is a packet-filtering device, such as a router or proxy server, that enables communication with only a designated application gateway within the internal network. Filtered external traffic is directed to this gateway, which examines it and redirects it to the appropriate destination.
- *Screened subnet gateway* – Also known as a *perimeter network* or *demilitarized zone* (DMZ), this is an additional subnet bounded by two firewalls that enclose a network segment containing components that should be accessible from both the Internet and the internal network, such as Web servers or e-mail servers. A variation of this topology includes a firewall containing three NICs that separate the internal network and the DMZ from the Internet.

In addition to these network topologies, you may have a need for an additional firewall or gateway within the internal network to protect network segments on which sensitive data is located, for example research or legal data. For detailed information on network segmentation design, refer to [Active Directory in Networks Segmented by Firewalls](#).

Designing a Service Recovery Process

An important, often overlooked, component of security design is planning for recovery of services that have failed or been compromised as a result of an incursion. One or more servers may have been damaged to the extent that they need partial or complete restoration, or you may have removed servers from the network to preserve evidence of the incursion for forensic analysis. The following are several components that a service recovery design or [contingency plan](#) should include:

- *A written recovery plan* – You shouldn't trust recovery procedures to memory, especially when a critical failure has suddenly occurred and each minute of downtime represents a large dollar loss to the company. The plan should include a prioritization of systems and services for recovery, a list of persons to be notified including legal authorities as required, special recovery procedures according to server roles (e-mail servers, firewall servers, IIS servers, DNS servers, and so on) and other system-specific procedures. You should test this plan thoroughly by recovering onto a duplicate network to ensure no components have been overlooked.
- *A proper backup and restoration design* – Make appropriate decisions as to the frequency and type of backups and the use of additional tools, such as [volume shadow copy](#) and [automated system recovery](#) (ASR). Furthermore, test your backups frequently by restoring them to alternate locations and verifying proper recovery.
- *Use of load balancing and clustering* – By spreading services across additional servers, you can remove and restore compromised machines without loss of service. In addition, these technologies offer improvements in service response at other times.
- *Use of hot, warm, or cold sites* – You can bring up your network at an alternate site should your main site go down. A hot site includes all components, such as servers and networks, to ensure a rapid recovery; however, note that warm and cold sites contain fewer components and need extra time to be brought online.

- *Recovery procedures for remote locations* – Servers at remote locations where IT staff are not present represent an additional challenge. Consider the use of technologies such as [Emergency Management Services](#) (EMS) as well as training the appropriate individuals as appropriate.
- *Audit the procedure and update it as required* – You should review processes, such as backup and restore procedures, and revise them as conditions change.

Technical Constraints in Security Design

Your security design will be limited according to the composition of your network and networks to which it must connect, the uses to which the network is put, required access from external locations, and other limitations.

Identifying Capabilities of the Existing Infrastructure

- Computer networks that have been built up over the years almost invariably contain older, legacy systems running operating systems, such as Windows 9x or Windows NT, that may not support the most recent security capabilities, such as Kerberos authentication or IPSec.
- At times it is impossible to eliminate computers running older operating systems because these computers serve a valuable role that is not easily updated; for example, a [TV station switching device or hospital equipment](#) running on Windows NT 4.0 computers.
- Microsoft Systems Management Server (SMS) provides a legacy client for Windows 98 and Windows NT 4.0. As discussed in [Legacy Client Security Environment](#), this client is not considered to be secure.
- By default, Windows 9x computers support only the very weak LAN Manager (LM) authentication protocol, and Windows NT computers support the first version of NT LAN Manager (NTLM). Although you cannot use Kerberos authentication on these systems, you can enable the use of NTLM v.2 by installing the Active Directory Client Extension from the Windows 2000 CD-ROM. To enable the use of NTLM v.2 on Windows NT 4.0 computers, install SP4 or higher.
- Remember that pre-Windows 2000 computers cannot apply group policy. You are limited to the capabilities provided with Windows NT 4.0 System Policy on these machines.
- Recognize any additional limitations that legacy clients may impose. For example, these clients may impose limitations on password strength. In addition, they will be unable to communicate with servers using the Secure Server (Require Security) IPSec policy.
- Consult [Hardening Microsoft Windows 98](#) for additional recommendations regarding legacy computer security.


Identifying Technology Limitations

- Hardware capabilities can limit your ability to upgrade an operating system. In some cases, it may be more feasible to replace outdated computers.
- If it is impossible to upgrade the operating system, which components of a security policy need to be modified or created exceptions to?
- If software applications in use preclude a hardware or operating system upgrade, can you accommodate this constraint within your security design?
- You may need to consider constraints accompanying different versions of infrastructure technologies on your network. Some of these technologies include e-mail and Web servers, naming servers such as DNS and Windows Internet Name Service (WINS), firewall and proxy servers, remote access servers, certificate servers, authentication infrastructure, customized software, and encryption (which can place a burden on server processing power).

Analyzing Interoperability Constraints

Networks that include different operating systems and that communicate with external networks of various types encounter various interoperability limitations. See [Enabling Interoperability with Kerberos Clients and Servers Running Other Operating Systems](#) for considerations that you should include in your security design.

- As well as computers running older Windows versions, many networks include systems such as UNIX, Linux, Macintosh, and even IBM-type mainframes. Analyze the constraints placed on network security by these systems, and accommodate them within your security design.
- Remote access scenarios, including telecommuters, branch offices, mobile sales staff, and foreign offices, all place limitations on a security design. For security considerations in planning a virtual private network (VPN), consult [Planning Security for a VPN](#).
- If your network must interface with networks operated by partner companies or clients, your security design must take the types of networks operated by these organizations into consideration. You may need different security designs for the interface with these networks, compared to access to and from other external locations.
- See [Web Services Interoperability](#) for references to security limitations imposed on Web services.
- We discuss additional constraints that limit security designs in later sections of this exam manual.



Ready to pass the 70-298 exam?

Download a **free** practice exam preview to find out if you're ready to pass.

Creating the Logical Design for Network Infrastructure Security

An appropriate logical design for a security infrastructure involves components that address authentication, confidentiality, integrity, and nonrepudiation of data being stored on and transmitted across the network. Without a solid logical design that is updated on a timely basis, network security is bound to fail at some point in time. This section looks at the use of Certificate Services, logical authentication strategies, network management security, and security update infrastructures.

Designing a Public Key Infrastructure using Certificate Services

A public key infrastructure (PKI) is the series of components designed to manage the issue, use, and maintenance of certificates on the network. It also includes certificate trusts and certificate revocation lists (CRLs). Consult [Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure](#) for comprehensive information on designing a PKI in Windows Server 2003.

Designing a Certification Hierarchy Implementation

The heart of the PKI is the certification authority (CA), which is a computer running Windows Server 2003 on which Certificate Services has been installed. You should be aware of the [four types of CAs](#) supported by Windows Server 2003:

- *Enterprise root CA* – Located at the top level, the root CA is the most trusted CA in any hierarchy of certificate servers. An enterprise root CA stores its database in Active Directory, thereby taking advantage of the security, fault tolerance, and replication mechanisms built into Active Directory.
- *Stand-alone root CA* – Also located at the top of its hierarchy, this server stores its database locally. While it is also the most trusted CA in its hierarchy, it is not associated with Active Directory.
- *Enterprise subordinate CA* – Located subordinate to a root CA, this server obtains a CA certificate from the root CA and stores its database in Active Directory. Being subordinate to another CA, it is never the most trusted CA in its hierarchy.
- *Stand-alone subordinate CA* – Also subordinate to another CA, it is not the most trusted CA in its hierarchy. Storing its certificate database locally, it is also not associated with Active Directory.

Although it is possible to deploy a single Windows Server 2003 computer running Certificate Services to issue certificates within your organization, it is common to use a [hierarchy of certificate servers](#). This enables you to keep the root CA offline, which protects it against compromise. Compromise of a root CA would render all certificates issued by any server in its hierarchy suspect and require the revocation of all certificates and the complete rebuilding of the hierarchy.

You can use either two- or three- tier certification hierarchies, as follows:

- A two-tier hierarchy (*Figure 2*) includes a root CA that issues certificates to one or more issuing CAs, which in turn issue all certificates required by the organization.




Figure 2 – A Two-tier CA Hierarchy Includes Two Levels of CA Servers

- A three-tier hierarchy (*Figure 3*) includes a root CA that issues certificates to intermediate CAs. These CAs issue certificates to issuing CAs, which in turn issue all certificates required by

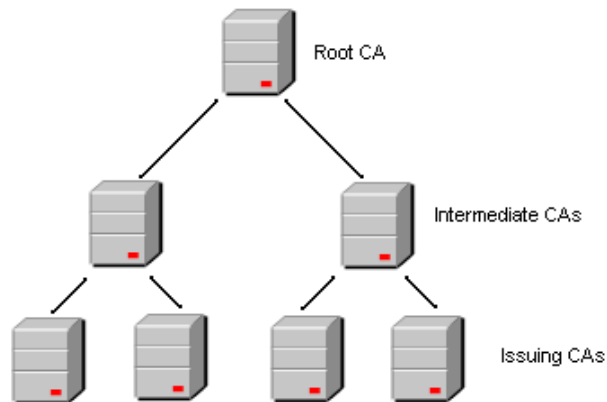


Figure 3 – A Three-tier CA Hierarchy Includes Three Levels of CA Servers

Several ways of designing a CA hierarchy are available, as follows:

- *Geographical* – Every subordinate CA (in either a two- or three-tier hierarchy) is responsible for issuing certificates to offices located in a specific geographic region. For example, you might have intermediate CAs situated in North America and Europe, and issuing CAs situated in the USA and Canada, and in England, France, and Germany.
- *Organizational* – Each subordinate CA issues certificates to different departments or divisions within a company. For example, different issuing CAs could be responsible for financial, design, legal, and sales departments.
- *Combination* – You can use a combination of these hierarchies; for example, a geographical

series of intermediate CAs each responsible for issuing CAs that issue certificates to various departments in its region.

- *Trusted* – You can set up trusts or cross-certifications that connect your company's PKI with that of a trusted partner or client company. We discuss cross-certification of Certificate Services later in this exam manual.

Designing Enrollment and Distribution Processes

After you have set up your CA hierarchy, you need to design the processes to be used for obtaining (enrolling) certificates and distributing them to the required computers and users. You have several means available for [certificate enrollment](#) and distribution:

- *Use of the Certificate Enrollment Web Pages* – When you install IIS and Certificate Services together on a Windows Server 2003 computer, the Certificate Enrollment Web Pages (*Figure 4*) are installed. These enable users with the proper permissions to request certificates for the required purposes.

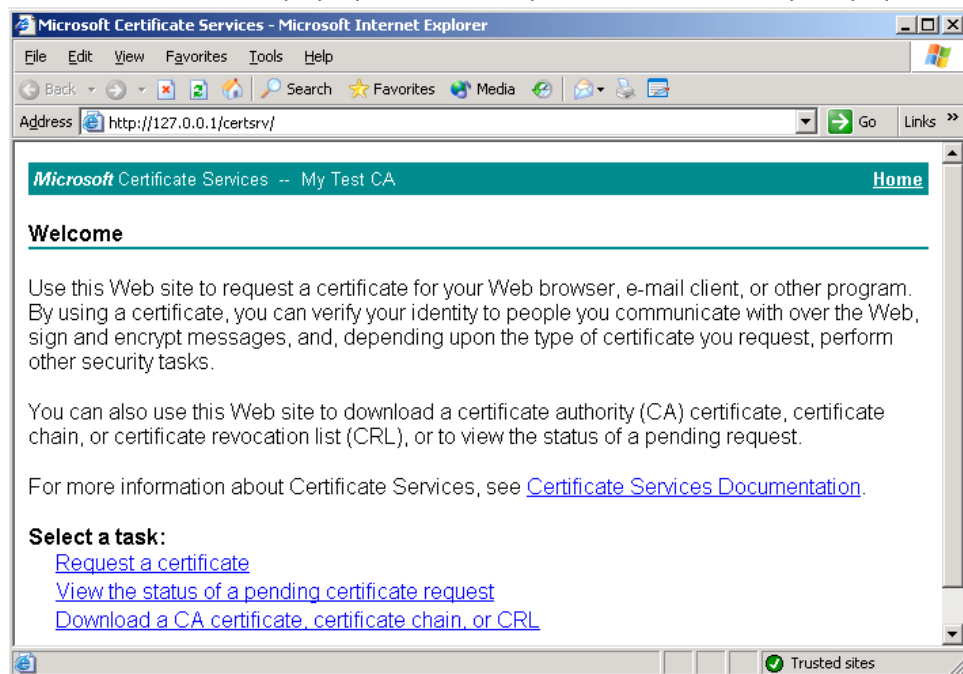


Figure 4 – The Certificate Enrollment Web Pages Enables You to Request Most Types of Certificates

- *Manual certificate enrollment* – Users or designated enrollment agents can use the Certificates snap-in to request certificates. This snap-in provides access to the [Certificate Request Wizard](#), which you can use to request a certificate for a computer, user, or service.
- *Use of the **Certreq.exe** command* – [This utility](#) enables you to create, submit, accept, and retrieve certificates, and write batch files or scripts to the automation of these actions. For more informa-

tion, type **certreq /?** at a command prompt from a Windows Server 2003 computer.

- *Autoenrollment of certificates* – Using this feature, which is new to Windows Server 2003, you can configure the [automatic enrollment](#) of certificates, including the retrieval of issued certificates and renewal of expiring certificates.
 - ▶ Autoenrollment is supported on certificates based on version 2 (Windows Server 2003) [certificate templates](#). However, you can duplicate a version 1 (Windows 2000) template to create a version 2 template.
 - ▶ The certificate server must be running Windows Server 2003 Enterprise Edition and client computers must be running Windows XP Professional or Windows Server 2003.
 - ▶ You can use Group Policy to configure certificate autoenrollment. The [Automatic Certificate Request Setup Wizard](#) assists you in configuring autoenrollment policies.

Issued certificates can be stored in one of several locations, including smart cards, files that can be exported to floppy disks, Active Directory, or intranet Web sites, among others.

Certificate Renewal

All certificates have an expiry date after which they are no longer valid. Consequently, you must design a [certificate renewal strategy](#).

- With time, encryption keys become more vulnerable to compromise. Periodic key renewal helps to reduce this risk.
- When the CA's certificate expires, all certificates issued by the CA, including subordinate CAs, also expire. You need to renew all associated certificates when you renew the CA's certificate. The expiry date set when certificates are issued or renewed can never be later than that of the CA's own certificate.
- To [renew the CA's certificate](#), right-click the CA in the Certification Authority snap-in and choose **All Tasks > Renew CA Certificate**. On the Renew CA Certificate dialog box, click **Yes** to create a new key pair or **No** to retain the current key pair, and then click **OK**.
- If you are using certificate autoenrollment, you can specify automatic certificate renewal.
- Available from the Certificates snap-in, the Certificate Renewal Wizard enables you to renew all certificates issued by your domain's CA.

Certificate Revocation

Certain events, such as the resignation of an employee or the compromise of a certificate or key pair, may require you to [revoke a certificate](#) before its expiry date.

- To revoke a certificate, right-click the certificate in the Issued Certificates node of the Certification Authority snap-in and choose **All Tasks > Revoke Certificate**. On the Certificate Revocation dialog box, select a reason code and then click **Yes**.
- Revoked certificates are published in the certificate revocation list (CRL). To publish a CRL, right-click Revoked Certificates and choose **All Tasks > Publish**. You can publish either of two CRL types:

- ▶ *Full CRL* – Includes all certificates revoked at any time by the CA.
- ▶ *Delta CRL* – Includes only those certificates revoked since the last publication of a CRL.
- To view the CRL or modify its properties, right-click Revoked Certificates and choose **Properties**. You can specify the publication intervals of full and delta CRLs and view CRLs in the Revoked Certificates Properties dialog box.

Certification Authority Auditing

The Certification Authority snap-in enables you to configure auditing of several events that are significant to operation of a CA. To configure auditing, you must enable the auditing of object access. You can then select the types of events to be audited from the Auditing tab of the CA's Properties dialog box, as shown in *Figure 5*.

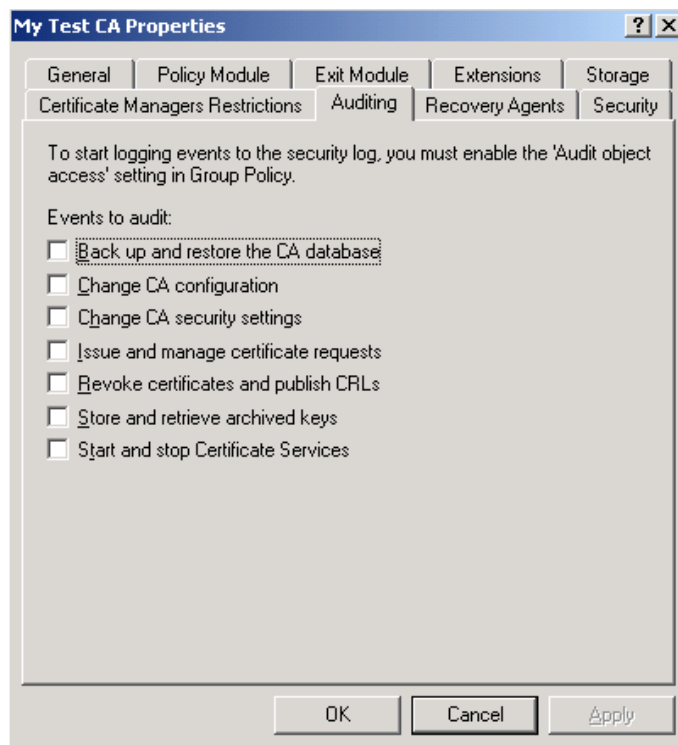


Figure 5 – The Auditing Tab of the CA's Properties Dialog Box

Designing Security for Certificate Servers

It is very important to harden certificate servers because of the roles they play. In particular, the compromise of a certificate server invalidates all certificates issued by the server and any subordinate CAs and requires that they all be revoked and reissued.

In addition to normal server hardening procedures, you can enable [role separation](#) on CA servers. Doing so enables you to specify that certain individuals are entitled to perform only a limited set of tasks on the CA server. The following roles are available:

- *CA Administrator* – Has the Manage CA permission, which enables the holder to configure the CA server, manage permissions, and renew CA certificates.
- *Certificate Manager* – Has the Issue and Manage Certificates permission, which enables the holder to manage certificate enrollment and revocation and initiate key recovery.
- *Auditor* – Has the Manage Auditing and Security Log permission, which enables the holder to configure, view, and maintain audit logs.
- *Backup Operator* – Can perform system backup and recovery, and can stop (but not restart) the certificate service.

You can use the Security tab of the CA server's Properties dialog box to assign these roles to the appropriate users or groups (Microsoft recommends that you assign them to groups). You can also use the [certutil](#) command to enforce role separation.

Designing a Logical Authentication Strategy

Designing Certificate Distribution

You can use Group Policy to manage the distribution of certificates on the network. Several policies are available in the [Public Key Policies](#) node of either Computer Configuration or User Configuration for use in your certificate distribution policy design.

- You can enable automatic certificate distribution by configuring a certificate autoenrollment policy.
- Users with Windows XP client computers automatically obtain certificates if you have configured autoenrollment using a version 2 template configured for automatic enrollment and the client has the Enroll permission on the template.
- You can also use a [registration authority \(RA\)](#) for Web-based enrollment. This is a server separate from the CA on which you have installed IIS and the Certificate Enrollment Web pages. Doing so enables you to provide improved protection for the CA server.
- When using an offline root CA, you should configure distribution of its certificate and CRL to appropriate network locations. The Extensions tab of the CA server's Properties dialog box enables you to configure distribution options.

Designing Forest and Domain Trust Models

Access to resources between domains or forests is dependent upon establishing the appropriate [trust relationship](#). By default, Active Directory in Windows 2000 or Windows Server 2003 provides complete two-way transitive trust relationships between all domains in a forest. You can also configure the following trust relationships to fulfill different needs and requirements:

- *External trust* – A trust relationship between two specific domains in different forests, or between an Active Directory domain and a Windows NT 4.0 domain or Kerberos realm. These are one-way non-transitive relationships.
- *Forest trust* – A trust relationship between all domains in each of two forests, both of which must be operating at the Windows Server 2003 forest functional level. These are transitive relationships, and can be one- or two-way.
- *Shortcut trust* – A trust relationship between two child domains in the same forest, this relationship serves to shorten the authentication path between the domains involved.

Trust relationships between forests carry potential security [threats](#), as follows:

- *Attack across a forest boundary by an intruder* – An intruder could monitor authentication data and obtain the security identification (SID) information for an administrative account. SID filtering, which is configured by default, is designed to help prevent this type of attack.
- *Attack on shared resources in a trusting forest by an intruder from a different forest* – The trust relationship enables an authentication path between the forests, which can create a larger attack surface. Use selective authentication, described below, to mitigate the potential of this attack.

When [designing a trust model](#), you need to take several factors into consideration:

- Determine the extent of access required. This will determine the type and direction of the trust relationship.
- Determine whether a new trust is really needed. Do currently configured trusts provide the appropriate access and security?

Determine the [level of authentication](#) required. You have the following options when creating the trust relationship with the New Trust Wizard:

- ▶ *Domain-wide authentication* – Automatically provides authentication for all users in the trusted domain accessing resources in the trusting domain. Users can access resources according to permissions, and have the access granted to the Everyone group in the trusting domain.
- ▶ *Forest-wide authentication* – Works similar to domain-wide authentication except that access is granted to all domains in the trusted forest. It is available only in forest trusts.
- ▶ [Selective authentication](#) – Enables you to limit the permission to authenticate to the appropriate groups or users. In Active Directory Users and Computers, allow the **Allowed to Authenticate** permission to the appropriate users or groups, and to the servers in the trusting domain to which they can be authenticated.

Designing Security that Meets Interoperability Requirements

Most modern networks contain several Windows versions, as well as various non-Microsoft computers such as UNIX, Linux, AS-400, and so on. It is necessary to design your authentication strategy to meet interoperability requirements with these systems. Designing the appropriate authentication strategy requires the knowledge of the following available [authentication protocols](#) and when each can be used:

- *LAN Manager (LM)* – Developed in the early days of Windows, this insecure protocol was used for authentication with Windows 3.x and 9x computers.
- *NTLM* – Used since the earliest versions of Windows NT, this protocol is still used when authenticating to computers running Windows NT prior to NT 4.0 Service Pack 4 (SP4).
- *NTLM v.2* – An improved version of NTLM that was introduced with Windows NT 4.0 SP4, and is used when authenticating to these computers. You can use NTLM v.2 with Windows 9x and older NT computers by installing the [Active Directory client](#) on these computers.
- *Kerberos* – The most secure, default authentication protocol of Active Directory in Windows 2000, Windows XP, and Windows Server 2003. You can also use this protocol for authenticating to UNIX and Linux computers located in a Kerberos V5 realm. Some AS-400 and mainframe installations support this authentication protocol. It provides single sign-on (SSO) capabilities for all network services.

You should restrict yourself to Kerberos authentication where possible, because this provides the most secure method of authentication. For details on enabling Kerberos authentication between Windows and UNIX environments, refer to [Microsoft Solution Manual for Windows Security and Directory Services for UNIX](#).

When pre-Windows 2000 computers are present, enable NTLM v.2. For more details on its use in Windows 9x and NT (pre-SP4) environments, refer to.

We look at client authentication strategies and the use of Group Policy for restricting the use of LAN Manager authentication later in this exam manual.

Establishing Account and Password Requirements for Security

Although many organizations are beginning to use advanced authentication technologies such as smart cards, tokens, and biometrics, the traditional user name and password is still the most common technology in current use. An understanding of how to design a password policy to balance security with user convenience is important.

The [password policy](#) provided by Group Policy includes the following settings:

- *Enforce password history* – Enable this policy and specify [at least 24 passwords](#) to be remembered. The user cannot reuse any of these passwords until the specified number of unique passwords have been used.
- *Maximum password age* – Enable this policy to force the user to change her password before the specified number of days have elapsed. In a moderate to high security environment, specify no more than 42 days.
- *Minimum password age* – Enable this policy to prevent the user from defeating the password history policy by cycling through a large number of passwords and returning to his favorite password. Specify a minimum age of at least one day.

- *Minimum password length* – Enable this policy and specify at least 7 characters in a moderate security environment and a higher number in a high security environment.
- *Passwords must meet complexity requirements* – Enable this policy to require that all passwords contain at least three of the four character types: uppercase letters, lowercase letters, numerals, and special characters.
- *Store passwords using reversible encryption* – Disable this policy for normal to high security. It actually weakens password security by storing passwords in a form available to certain non-Microsoft implementations, such as Macintosh, or when using Challenge Handshake Authentication Protocol (CHAP) or digest authentication in IIS.

In addition, several user account properties affect passwords. You can select the **User Must Change Password at Next Logon** option to force a user to change his password. For service accounts whose password should not change, you can select the **User Cannot Change Password** and **Password Never Expires** options. You should never select these options for regular user accounts.

Together with password policies, configure an [account lockout policy](#) that locks a user out after the specified number of incorrect passwords have been entered. Set this policy to a high enough value (say 10 or more) so that users mistyping their passwords do not inadvertently lock themselves out, causing a help desk call. Password cracking attempts usually run through a very large number of passwords and will still cause lockout at any reasonable number of attempts.

In conjunction with the maximum password length policy, you should enable the **Interactive logon: Prompt user to change password before expiration** policy, which warns the user her password is about to expire. Set this policy to the number of days' notice you want the users to receive. Along with several other policies affecting logon, it is found at the [Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options](#) node of Group Policy.

Designing Security for Network Management

Managing the Risk of Managing Networks

Whatever means you employ for managing your network, risks are involved. No matter how trustworthy an administrator is, temptation can set in and he might decide to perform actions such as transfer of funds out of a corporate account, or copy confidential information such as research or legal records to a USB flash drive, or so on. No administrator is infallible and some might inadvertently create a security risk when performing some legitimate action. Or they might decide to ease up on security because network performance has declined.

When designing security for a large corporate network, you should be aware of the concepts of autonomy and isolation and when to use each:

- *Autonomy* means that although a segment of a network is managed by local administrators, external control of the network is still possible. Examples of autonomy include OUs or child domains.

Isolation means that a precise, absolute security boundary exists between two segments of the network. This could simply involve a stand-alone server that does not belong to the domain and therefore cannot be administered by domain administrators, or it could imply a separate forest. In particular, you should note that a domain is not an absolute security boundary (unlike the Windows NT 4.0 domain) because all domains in a forest trust each other by default and forest-wide groups such as Enterprise Admins and Schema Admins exist. For a case study on the use of isolation, refer to [Server and Domain Isolation Using IPSec and Group Policy](#).

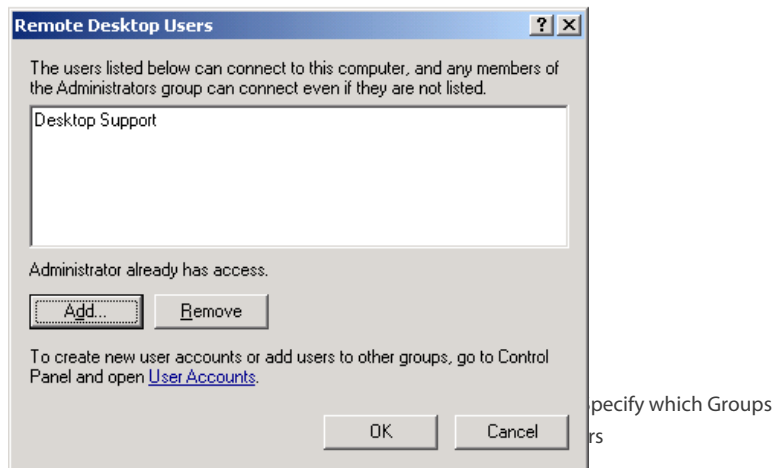
The following are several best practices you should follow for reducing the risk of network management:

- Determine the need for, and extent of, autonomy or isolation. This includes involvement of corporate managers and executives to ensure that overall company security needs are met.
- Limit the number of individuals assigned to groups such as Domain Admins, Enterprise Admins, and Schema Admins. Use the [Restricted Groups](#) policy in Group Policy where feasible.
- Use the principle of least privilege when assigning administrative control. In other words, provide only the minimum amount of authority to individuals and groups that will be managing portions of the network. See [Best Practices for Delegating Active Directory Administration](#) for additional details.
- The [Delegation of Control Wizard](#) facilitates assigning the appropriate extent of control to users and groups charged with administering portions of the network.
- Use auditing to detect misuse of any administrative privileges by the various groups. We discuss auditing later in this exam manual.

Designing Server Administration by Using Common Administrative Tools

Microsoft provides several administrative tools, such as Microsoft Management Console (MMC), Terminal Services, Remote Desktop for Administration, and Remote Assistance, for server administration purposes. The following are several best practices you should follow in designing server administration with the use of these tools:

- Place administrative users into specific groups and delegate only the required extent of authority to each group.
- Create customized MMC consoles designed for the administrative groups you have created. To ensure that consoles are not modified to add increased functions, configure them for a specific mode such as user mode—limited access, multiple windows, or user mode—limited access, single window. See [Console access options](#) for more information.
- Provide administrative users with two user accounts, one for administrative activities and a second one for general use such as e-mail and Web access. Instruct these users to always log on with the general account and use the **Runas** option to access the administrative account as required.
- On servers running Windows Server 2003, it is not necessary to install Terminal Server if you only need to use [Remote Desktop for Administration](#). Only install Terminal Server if users need to connect to the server for running applications remotely.
- If administrative users need to access servers remotely, add the groups containing these users to the Remote Desktop Users group. On the Remote tab of the System Properties dialog box, click **Select Remote Users**. On the Remote Desktop Users dialog box (*Figure 6*), click **Add** to add the appropriate groups.
- Remote Assistance is designed more for users requiring assistance or tutoring in the use of specific computer functions than for the remote administration of servers. For additional information on using Remote Assistance, refer to [Using Windows Server 2003 in a Managed Environment](#).



Designing Security for Emergency Management Services

[Emergency Management Services](#) (EMS) is a new feature in Windows Server 2003 that is designed to enable the administration of a remote server without the availability of normal network drivers and services (also known as out-of-band administration). This includes the ability to restart a non-responsive server. You can perform tasks such as viewing Stop errors, restarting the server, viewing power on self-test (POST) messages, using Remote Installation Services (RIS) to install the operating system, and all the normal server management tasks. EMS includes the [Special Administration Console](#) (SAC), which is a command-line tool from which you can perform administrative tasks. If the server is not responding properly, you can administer it by using a special version of SAC known as ISAC.

EMS represents a potential attack vector for unauthorized access to your servers. The following are several [best practices](#) you should follow in securing EMS:

- Use in-band management tools such as Remote Desktop for Administration when servers are performing normally.
- Use a separate management network to connect servers. The servers are connected to this network by means of hardware devices known as *terminal concentrators*. In particular, select a terminal concentrator that provides security features for network connections.
- Use secured communication methods such as Telnet over Internet Protocol Security (IPSec) or Secure Shell (SSH). In addition, follow other network security practices such as using strong passwords and encrypting network communications.
- Limit physical access to the servers, terminal concentrators, and other networking devices.
- If using a modem for connections with EMS, use callback to ensure connections to the network originate only from authorized locations.

Designing a Security Update Infrastructure

Microsoft releases security updates of various types on a fixed schedule, normally the second Tuesday of every month. They may also release urgent updates at other times as required. The common types of security updates are as follows:

- A [patch](#) (can also be called a security update) is an update that addresses one or a small number of vulnerabilities or configuration problems.
- Microsoft considers a *hotfix* to be a single, cumulative package that addresses a specific problem that has been identified by a company through Microsoft Product Support Services.
- A *service pack* is a cumulative series of hotfixes, security updates, and other updates. It also might include new functionality, such as the Security Center, incorporated into Windows XP SP2 and Windows Server 2003 SP1.

Designing a Software Update Services (SUS) Infrastructure

SUS (to be replaced in 2005 by Windows Server Update Services—WSUS) provides a centralized location for deploying all Microsoft security updates to computers on your network. It provides the following advantages:

- You can set up a parent-child SUS server architecture that supports up to 15,000 clients per SUS server. You can copy updates from a primary SUS server to secondary SUS servers without the need for the secondary servers to be connected to the Internet.
- You can approve individual updates on every SUS server. Doing so allows you to test and deploy updates on a scheduled basis.
- You can configure client computers to obtain their updates from the SUS server. This also includes computers that are not connected to the Internet.

The following are several guidelines for use when designing a SUS infrastructure (see [Solution Guidance](#) and [Operational Guidance](#) for details on designing and operating SUS):

- Deploy a test network containing computers representative of all types used on the network and use this network to test all updates before approving them for use on the production network. Virtualization products such as Microsoft Virtual PC or Virtual Server are useful in reducing the amount of hardware required for the test network.
- Set up a hierarchy of SUS servers so that the parent SUS server has a high-bandwidth Internet connection available and the child SUS servers are located close to all client computers (including servers) that they will serve. Carefully document all test results in case problems are encountered later on the production network.
- Another reason to use multiple child SUS servers is the need to approve certain updates only to portions of the production network (for example, in case an update affects the functionality of an application that is in use in a single department only).
- Deploy a [Network Load Balancing](#) (NLB) cluster of child SUS servers if you must deploy patches to a large number of computers. Such a design provides redundancy and load balancing, ensuring prompt delivery of updates to all client computers.

- Ensure that all SUS servers are protected against intrusion. Compromise of SUS servers could result in compromise of all network computers. In addition, ensure that SUS servers are protected with the latest antivirus signatures.
- Don't forget to include a strategy for updating portable computers and other machines that are not continuously connected to the network in your design.
- Design an appropriate backup strategy for the SUS servers, including the SUS Web site, its content folder, and the IIS metabase.

Using Group Policy to Enable Deployment of Software Updates

Group Policy provides several options that you can use to deploy software updates to client computers. To access these policies, open the Group Policy Object Editor focused on a GPO linked to the appropriate Active Directory container, navigate to **Computer Configuration\Administrative Templates\Windows Components\Windows Update**, and configure the following policies:

- *Configure automatic updates* – Enable this policy and select the method to be used for automatic updating (see Figure 7):
 - ▶ *Notify for download and notify for install* – Prompts the user to download and install updates.
 - ▶ *Auto download and notify for install* – Downloads updates and notifies the user when they are ready for installation.
 - ▶ *Auto download and schedule the install* – Automatically downloads and installs updates on the days and times you specify in the drop-down lists provided. This is the recommended option.

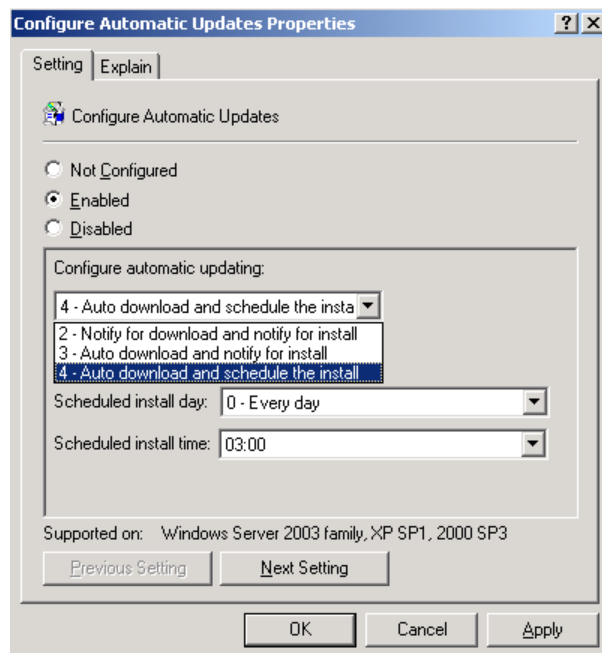


Figure 7 – Group Policy Provides Several Options for Configuring Automatic Updates

- *Specify intranet Microsoft update service location* – Enable this policy and specify the name of the SUS server to be used for downloading updates. On a large network with multiple child SUS servers in different locations, configure GPOs linked to the appropriate sites, domains, or OUs with their specific child SUS servers.
- *Reschedule Automatic Updates scheduled installations* – Enable this policy to ensure that computers that are not online at the time of download receive their updates, and specify the number of minutes to wait after startup.
- *No auto-restart for scheduled Automatic Updates installations* – Prevents computers from automatically restarting when an update that requires a reboot is installed. You should restart affected computers manually at a convenient time if you have enabled this policy.

For additional information on designing a Group Policy for use with SUS, refer to [Deploying Patches with Software Update Services 1.0](#).

Identifying Computers that are not at the Current Patch Level

After you configure SUS to deploy patches and apply a policy in Group Policy, verify the installation of patches on the computers on your network, and then take steps to refine your patching design should you determine that patches have not been installed as required. You may even want to perform this verification before you configure SUS to assess the state of your patch deployment.

Microsoft Baseline Security Analyzer (MBSA), shown in *Figure 8*, scans computers running Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003 for security vulnerabilities. It checks whether current patches have been installed, and determines the need for further system hardening.

The screenshot displays the Microsoft Baseline Security Analyzer (MBSA) interface. The main window shows a 'View security report' for the computer 'CRAMSESSION\SERVER1'. The report includes the following details:

- Computer name:** CRAMSESSION\SERVER1
- IP address:** 192.168.1.1
- Security report name:** CRAMSESSION - SERVER1 (3-4-2005 11-46 AM)
- Scan date:** 3/4/2005 11:46 AM
- Scanned with MBSA version:** 1.2.4013.0
- Security update database version:** Security updates scan not performed
- Security assessment:** Severe Risk (One or more critical checks failed.)

The 'Windows Scan Results' section lists the following vulnerabilities:

| Score | Issue | Result |
|-------|-----------------------------|--|
| ✗ | Local Account Password Test | Some user accounts (1 of 6) have blank or simple passwords, or could not be analyzed. What was scanned Result details How to correct this |
| ✗ | Password Expiration | Some user accounts (2 of 8) have non-expiring passwords. What was scanned Result details How to correct this |
| ℹ | Windows Firewall | Windows Firewall is disabled or has exceptions on all network connections. What was scanned Result details How to correct this |
| ✓ | File System | All hard drives (1) are using the NTFS file system. |

At the bottom of the report, there are buttons for 'Previous security report' and 'Next security report'.

Figure 8 – MBSA Checks for a Large Range of Vulnerabilities on Windows 2000 and Later Computers

The following are the major vulnerabilities that MBSA checks for:

- *Windows vulnerabilities* – Windows Firewall settings (Windows XP SP2, Windows Server 2003 SP1, and later), weak or blank passwords, missing patches or other updates, unnecessary services running, local administrators, and so on.
- *IIS and SQL vulnerabilities* – On servers running either of these applications, MBSA checks for vulnerabilities that are specific to these applications.
- *Additional security configurations* – MBSA checks for Internet Explorer and Outlook security zone settings, Microsoft Office macro settings, and Windows Media Player vulnerabilities.

You can use the **mbsacli.exe** utility to run MBSA from a command line. This enables you to script the execution of MBSA across multiple computers, and to schedule its execution using the Task Scheduler. When using MBSA together with SUS, you can use the GUI to point MBSA to the SUS server, or from the command line as follows:

```
mbsacli /d domain_name /SUS http://susserver
```

In this command, all computers in the specified domain are checked against approved updates from the specified SUS server.

For additional information on using MBSA from the command line, refer to [How to Use Microsoft Baseline Security Analyzer \(MBSA\)](#).

Creating the Physical Design for Network Infrastructure Security

The requirements for securing your physical network are wide ranging and encompass all components of the network infrastructure, including such items as wireless networks, virtual private networks (VPNs), communications with customer and partner organizations, and Web communications. This section looks at creating designs for all facets of physical network security.

Designing Network Infrastructure Security

Specifying Firewall Protocols

The firewall is the border of your network. Think of it in the same way as a country thinks of its borders with neighboring countries; it uses border guards and controls to secure these borders. In the same way, you specify border controls for your network at the firewall. Choose the firewall that's appropriate for your organization and then design a policy that specifies the types of traffic that will be permitted across the firewall. Refer to [Enterprise Design for Firewalls](#) and [Perimeter Firewall Design](#) for details on all facets of firewall design.

An overall firewall design includes the following components:

- *Logical design* – This phase involves defining the appropriate firewall functions. Possible designs include single-tier, two-tier, and multiple-tier design. See [Enterprise Design for Firewalls](#) for more information on these designs.
- *Device design* – This phase involves selection of the appropriate type of firewall for the network.

Possible firewall designs include packet filtering, network address translation (NAT), circuit-level inspection, proxy servers, and stateful inspection filters.

- *Physical design* – At this point, you must develop the physical hardware and software configuration of the internal network, demilitarized zone (DMZ), and firewall location.

Keep in mind the following considerations when designing your firewall configuration:

- It is best to start by blocking all traffic by default and then configure exceptions according to the traffic that must cross the firewall. Microsoft [Internet Security and Acceleration \(ISA\) Server](#) operates on this premise.
- Carefully consider which ports should be opened. You may be allowing more traffic than you think. Set up a formal approval process that ensures only the required ports and protocols are allowed through the firewall.
- An Application layer firewall examines traffic to detect malicious traffic that uses a common port such as port 80, while allowing normal traffic such as requests to and replies from your Internet Web server.
- Proxy servers, such as ISA Server, can act as border controls, requiring authentication and restricting resource access. You can impose limits such as the IP addresses that are allowed, or file types that can be downloaded. Most proxy servers create their own connections and relay them to the requesting client, rather than establishing a direct connection from the client to the resource.
- Required ports and protocols depend on the services you must allow through the firewall. See [Configuring Firewalls](#) for information on the use of firewalls with servers running Microsoft Exchange Server.

Specifying IP Packet Filtering

TCP/IP filtering is a simple but powerful means of controlling access to your servers. It limits access to the server by enabling you to specify the TCP ports, UDP ports, and IP protocols allowed to access the server. Perform the following procedure to configure IP filtering:

1. Access the **Properties** dialog box for the network adapter for which you want to configure IP filtering.
2. Select **Internet Protocol (TCP/IP)** and click **Properties**.
3. On the Internet Protocol (TCP/IP) Properties dialog box, click **Advanced**.
4. Select the **Options** tab, ensure that **TCP/IP filtering** is highlighted, and then click **Properties**.
5. On the TCP/IP Filtering dialog box (see *Figure 9*), select **Enable TCP/IP Filtering (All adapters)**.
6. As required, select **Permit Only** under one or more of TCP Ports, UDP Ports, and IP Protocols, and then click **Add**.
7. On the **Add Filter** dialog box, type the port or protocol number that should be allowed, and then click **OK**. Repeat as required to add all required ports or protocols, and then click **OK** to close the TCP/IP Filtering dialog box.

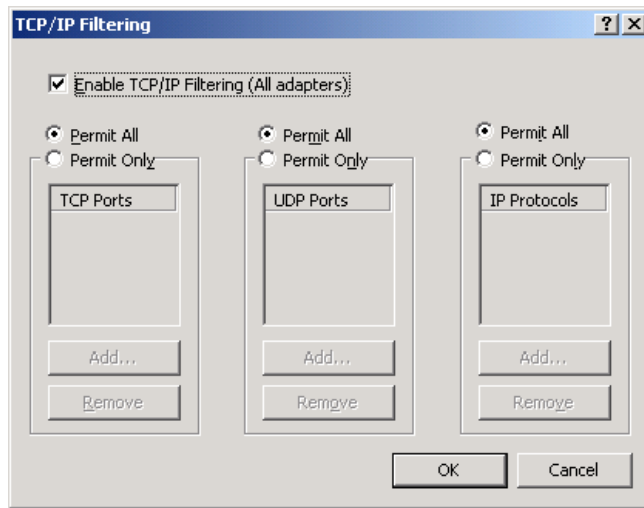


Figure 9 – The TCP/IP Filtering Dialog Box Enables You to Restrict the Ports and Protocols that Are Allowed to Access Your Server

Designing Internet Protocol Security (IPSec) Policies

Group Policy in Windows Server 2003 provides default IPSec policies that you can assign in GPOs linked to an Active Directory site, domain, or OU. Although it is possible to use one of the default IPSec policies, you can create customized IPSec policies that provide the appropriate level of security for your organization's needs. See [Designing an IPSec Policy](#) for additional information on the IPSec security policy design process.

At any time you can have only one IPSec policy assigned in a given GPO. An IPSec policy is made up of the following components (see [How IPSec Works](#) for detailed information and for additional information on creating policy rules):

- *One or more IP Security rules* – You can invoke any number of rules by selecting them from the Rules tab of the policy's Properties dialog box.
- *Filter lists* – Specified on the IP Filter list tab of the rule's Edit Rule Properties dialog box, you can select only a single filter list per rule. Each filter list can contain any number of IP filters that specify the source and destination ports, protocols, IP addresses, and subnet masks.
- *Filter actions* – Specified on the Filter Action tab of the rule's Edit Rule Properties dialog box (*Figure 10*), the filter action determines the type of action that will be undertaken for traffic matching any one of the rule's filter lists. Only one filter action is permitted for a given rule. The action specifies that traffic meeting the filter will be permitted or blocked, or security will be negotiated according to additional criteria that can be specified on the Security Methods tab of the action's Properties dialog box. See [Filter Action](#) for more information.

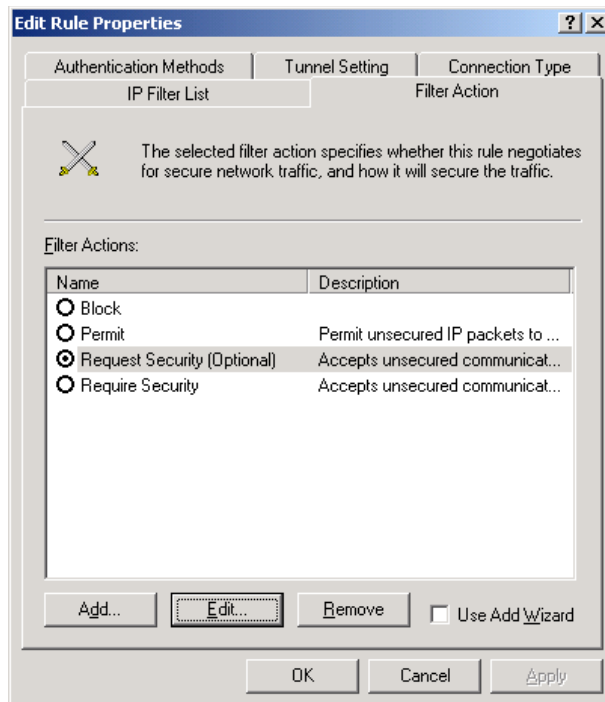


Figure 10 – You Can Configure the Filter Action Used by an IPSec Rule from the Filter Action Tab of the Rule's Edit Rule Properties Dialog Box

- *Additional security method settings* – When you have selected Negotiate Security as the filter action, you can configure the following security methods:
 - ▶ *Integrity and encryption* – Uses Encapsulating Security Payload (ESP) to encrypt data in transit and verify that data is not modified.
 - ▶ *Integrity only* – Uses Authentication Header (AH) to verify that data is not modified in transit. Data is not encrypted.
 - ▶ *Custom* – Enables you to configure additional integrity and encryption settings (*Figure 11*) including the integrity algorithms (MD5 or SHA1), encryption algorithms, such as Data Encryption Standard (DES) or Triple DES (3DES), and session key settings (which determine how frequently new encryption keys are generated).

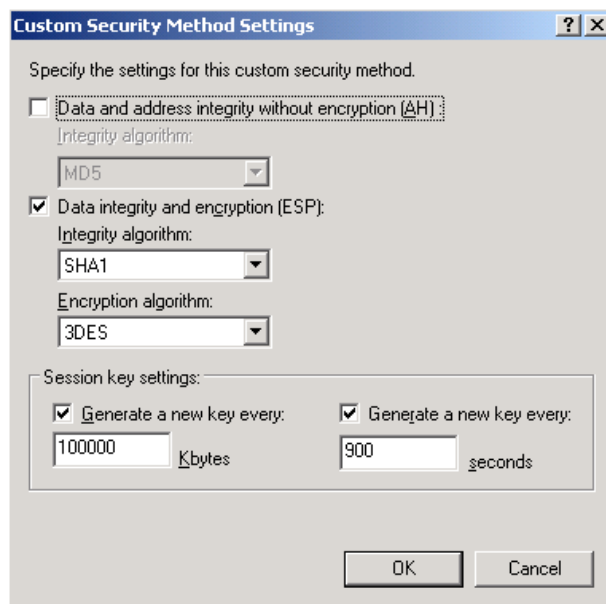


Figure 11 – The Custom Security Method Settings Dialog Box Enables You to Specify Additional Security Settings Related to Security Negotiation

The following are several guidelines you should consider when [designing an IPSec policy](#):

- *Fit the policy design to your Active Directory structure* – IPSec policies are stored in the IP Security Policies container, which is separate from the GPOs that apply the policies. Ensure that the proper permissions are applied to this container so that policies apply as expected.
- *Consider the impact of routers and firewalls* – Ensure that routers and firewalls are configured to pass IPSec-secured traffic.
- *Understand the capabilities of computers and other devices on the network* – Older devices may be unable to communicate if IPSec is required.
- *Consider the impact of securing data on network processing resources* – Using IPSec to encrypt all network traffic may cause an undesirable load on processing power. Certain types of network traffic, such as publicly available data, do not need to be secured.
- *Use IPSec-enabled network adapter cards on servers* – These adapters can encrypt and decrypt data on-the-fly in hardware, thereby removing this load from server processors.
- *Do not encrypt authentication traffic to domain controllers* – IPSec cannot be negotiated before a client is authenticated, so the client must be able to connect to the server before IPSec is set up.
- *Test your IPSec design thoroughly before deploying it* – Successful implementation of IPSec can be tricky and can cause unforeseen communication problems. Deploy the IPSec policies on a test network with comprehensive logging and ensure that vital communications are not blocked.

Securing a Domain Name Service(DNS) Implementation

An Active Directory network requires DNS to locate resources including domain controllers and objects referenced in the Active Directory database. Compromise of DNS could enable an intruder to redirect Active Directory queries to spoofed servers thereby providing information that enables the intruder to attack sensitive resources. Permitting unsecured zone transfers enables an intruder to obtain an entire copy of your DNS zone data, providing a footprint of the network. A denial of service (DoS) attack on DNS could bring all operations to a halt. See [Security Information for DNS](#) for more details on threats to DNS and mitigation procedures.

The following are several recommendations for securing DNS on your network:

- *Ensure the appropriate placement of DNS servers* – DNS servers that resolve names of internal network computers should be located on the internal network. You can configure forwarding of requests for external network addresses to servers on the DMZ or hosted by your Internet service provider (ISP).
- *Use split DNS* – Configure a separate DNS subdomain for internal addressing with a parent DNS domain located in the DMZ that provides name resolution for external connections to publicly available resources such as Web and e-mail servers.
- *Use Active Directory-integrated DNS zones* – these are automatically encrypted and replicated with other Active Directory information. Furthermore, you can configure the [Secure only dynamic update option](#) to ensure that only authenticated clients can update DNS zone information.
- *Restrict zone transfers* – Use the [Zone Transfers tab](#) of the zone's Properties dialog box (*Figure 12*) to specify which servers are allowed to receive zone transfers. This is especially valuable if you are using zones that are not Active Directory-integrated.

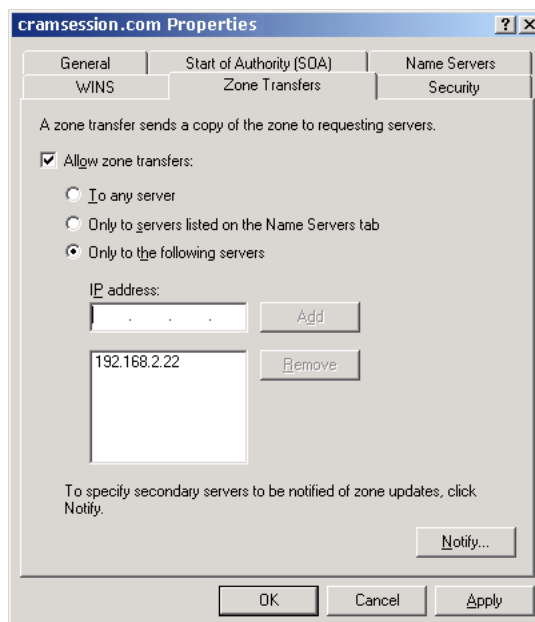


Figure 12 – Use the Zone Transfers Tab to Limit the Servers that Receive Zone Transfers

- *Ensure that DHCP servers are secured* – If you have configured Dynamic Host Configuration Protocol (DHCP) to provide DNS server information to clients, compromise of the DHCP server could enable an attacker to insert invalid DNS information that redirects clients to the attacker's servers or result in a DoS condition.
- *Ensure that the DNS cache is secured against pollution* – Found on the Advanced tab of the server's Properties dialog box, this setting prevents the addition of bogus data to the server's cache, which could redirect clients to an attacker's server.
- *Encrypt replication traffic* – When replication traffic must cross a wide area network (WAN) connection, set up an IPSec tunnel, which authenticates each server to the other before replication can take place. This action also provides a VPN-tunneled connection that can cross unsecured connections such as the Internet.
- Use [DNSSEC security extensions](#) – These enable data origin authentication and integrity checking, thereby providing further protection for DNS zone data.

Designing Security for Data Transmission

Security for data transmission can involve any of several protocols according to the type of network across which the data is transmitted:

- Use IPSec to secure data transmitted across local networks and VPNs. We have already discussed the design of IPSec policies.
- Use Secure Sockets Layer (SSL) to secure traffic to and from Web servers. We discuss the use of SSL later in this exam manual.
- Use VPN technology, including either Point-to-Point Tunneling Protocol (PPTP) encryption or Layer Two Tunneling Protocol (L2TP) together with IPSec to create a secure tunneled connection across the Internet between two remote locations. We discuss VPN security later in this exam manual.

Designing Security for Internet Information Services

IIS 6.0 in Windows Server 2003 includes applications for several [Internet-based communications](#), including the World Wide Web (WWW), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and Network News Transfer Protocol (NNTP). These applications are all targets for attack because of their intimate relationship with the Internet and Web-based business.

For a comprehensive review of IIS hardening methodologies, consult [Managing a Secure IIS 6.0 Solution](#) and references cited therein.

Designing IIS User Authentication

IIS provides the following [five methods](#) for authenticating users visiting Web sites (see *Figure 13*). In addition to these methods, you can use certificates or Remote Access Dial-In User Service (RADIUS) for authentication:

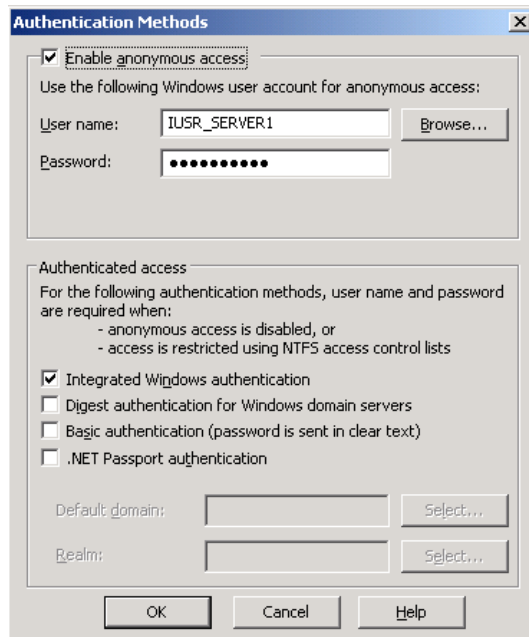


Figure 13 – The Authentication Methods Dialog Box Enables You to Select Different Means for Authenticating Visitors to Your Web Site

- *Anonymous access* – Enables visitors to access public Web sites without a user name or password. Visitors are authenticated using the default IUSR_servername account.
- *Integrated Windows authentication* – Uses Windows domain user accounts to authenticate visitors by means of the NT LAN Manager (NTLM) or Kerberos protocols. Requires Internet Explorer 4 or later.
- *Digest authentication for Windows domain servers* – Uses a Message Digest 5 (MD5) hash to protect passwords of users stored in Active Directory. Users must have domain user accounts and use Internet Explorer 5 or later to connect to the Web server.
- *Basic authentication* – Sends credentials in clear text, but allows users to authenticate from any type of browser. You can use SSL or IPSec encryption mechanisms to protect against password sniffing.
- [.NET Passport authentication](#) – Uses SSL to send credentials to a .NET Passport Web service, which returns a cookie containing a valid access ticket. You do not need to maintain your own user account database, and users can use the same name and password to access multiple Web sites.

Use of RADIUS for IIS Authentication

When using dial-up or VPN remote access to an IIS server for purposes such as Web site administration or access to mailboxes by means of [Outlook Web Access](#), you can use a RADIUS server, such as Microsoft Internet Authentication Service (IAS), to authenticate users. We discuss VPN access and authentication later in this section.

Use of Certificates for IIS Authentication

IIS 6.0 also supports authentication by means of user (client) certificates. Used in conjunction with basic, digest, or integrated Windows authentication, certificates provide a high level of protection against unauthorized access.

To use client certificates with IIS, you must first obtain and install a [server certificate](#) for the IIS server.

On the Directory Security tab of the Web Site Properties dialog box, you can configure [certificate-based authentication](#) by clicking **Edit** under **Secure Communications**. This displays the Secure Communications dialog box, from which you can configure the following options:

- *Require secure channel (SSL)* – Requires that all connections be made using SSL and a URL that begins with https:.
- *Client certificates* – You can choose to ignore, accept, or require client certificates. You must select the **Require Secure Channel** option to require client certificates.
- *Enable client certificate mapping* – Maps client certificates to Windows user accounts. See [Mapping Strategies](#) for suggestions for using certificate mapping. You can choose either of the following mapping methods:
 - *One-to-one* – Maps individual certificates to Windows user accounts.
 - *Many-to-one* – Uses wildcard mapping rules to accept any certificate that meets the rules you have specified.
- *Enable [certificate trust list](#)* – This is a list of root CAs whose certificates are trusted for authentication. Enabling this option allows you to trust only certificates issued by your company's and partner company's CAs.

Designing IIS Security Baselines According to Business Needs

We introduced the need for analyzing business requirements as they relate to security at the beginning of this exam manual. Microsoft has improved the security of IIS 6.0 relative to earlier versions by not including it in default installations of Windows Server 2003 and installing it in a locked down condition that serves only static Web content when first installed from the Windows Components Wizard.

Designing security for IIS involves collaboration with managers and others to ensure that you achieve the proper balance between functionality and security. Ensure that the managers understand that the longer the Web sites are unavailable because of security breaches, the greater the financial loss to the company; furthermore, you should attempt to provide an estimate of potential losses. See [Security in IIS 6.0](#) and links contained therein for security design recommendations. You can find additional recommendations in [IIS 6.0 Security Best Practices](#).

The following are several considerations you should address in your IIS security baseline design:

- *What types of content will be served by the IIS server?* This determines whether you need to install additional IIS components, such as FTP, SMTP, and NNTP, as well as services such as Active Server Pages and Web Distributed Authoring and Versioning (WebDAV). For a review of available IIS components, refer to [How to Identify IIS 6.0 Components in Windows Server 2003](#).
- *What level of access do different groups need to data on the IIS server?* This includes anonymous guests, authenticated users, Web site developers, and administrators. You can use a balance of [IIS Web site permissions](#) and NTFS permissions to provide the appropriate levels of access. Windows Server 2003 provides a new tool called [Authorization Manager](#), which provides additional management of access control in IIS 6.0.
- *What type of authentication is most suitable for visitors to the Web sites?* We have already reviewed the available authentication types.
- *Where will the IIS server be located?* Although a public Web server is not often located on the perimeter network, you can also outsource it to your ISP, who then takes on the responsibility of keeping it secure. An intranet Web server only accessible to employees is normally placed on the internal network.
- *How will traffic to and from the IIS server be secured?* Do you need to implement SSL for external communications? If the IIS server needs to communicate with internal servers, such as SQL servers, consider using IPSec filters. Refer to the Windows Server 2003 Security Manual discussed later in this exam manual for recommendations on IPSec filters for IIS servers.
- *How much auditing and logging is needed?* We discuss designing IIS monitoring later in this section.

Designing Web Site Security By Enabling the Minimum Required Services

As shown in *Figure 14*, IIS is a component of Application Server, which you can install from the Windows Components Wizard. In turn, IIS includes several components, some of which contain subcomponents, as shown for the World Wide Web service. Consider the types of uses required of the server to decide which subcomponents are required. Refer to [Enabling Only Essential IIS Components and Services](#) for a description of the available services and subcomponents, as well as suggestions for several types of IIS installations. Also see [Enabling and Disabling Dynamic Content in IIS 6.0](#) for information on configuring dynamic content.

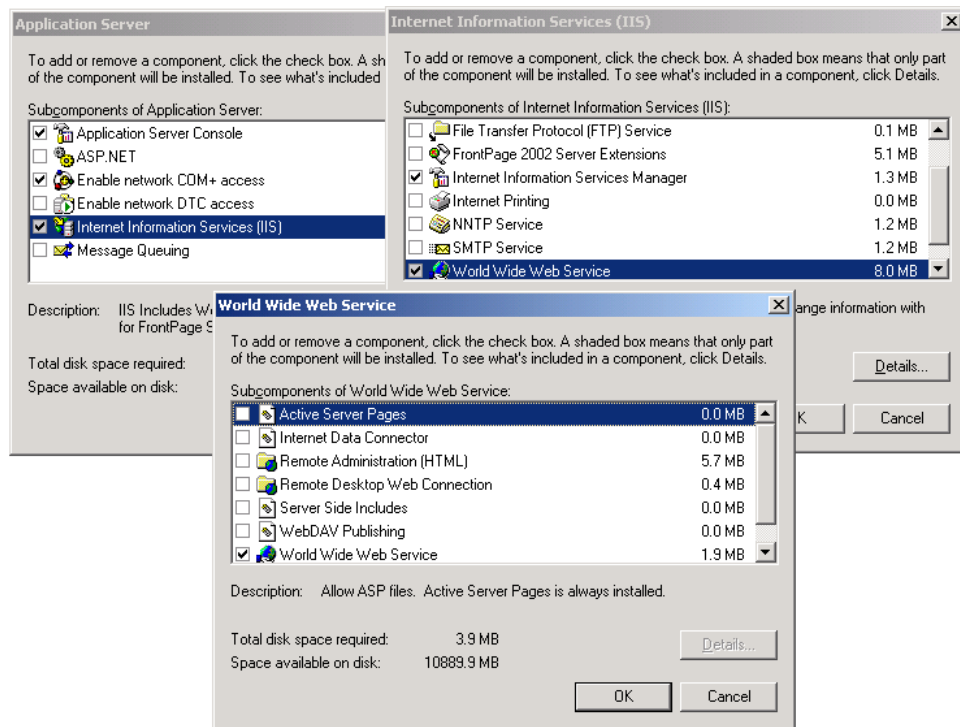


Figure 14 – Installing IIS 6.0 from the Available Components of Application Server in the Windows Components Wizard

In addition to these components, you should review the Windows services running on your Web server and disable any that are not explicitly required. Consider disabling the following services.

- *Remote Registry* – Used for remote management of the server. On a high-security Internet server, you should disable this service and use an alternative method for managing the server.
- *Terminal Services* – You should not allow ordinary users to run terminal sessions with a Web server. The administrator can still use Terminal Services Administration mode to administer the Web server.
- *SMTP* – If you enable SMTP, ensure that uncontrolled relay is not permitted, or intruders could use your server as a gateway to send spam.
- *Isolation Mode* – If you are running applications that were originally developed for IIS 5.0 in Windows 2000, you may need to run your server in [IIS 5.0 isolation mode](#) for the application to run properly.

Designing a Monitoring Strategy for IIS

An appropriate [logging](#) and monitoring strategy must be in place for your IIS servers to track access to Web sites, files, and applications. As well as informing you which resources are most popular with visitors to your sites, such a strategy provides information on successful and failed intrusion attempts on your servers.

The Web Site tab of the Web Sites Properties dialog box (or of the Properties dialog box for each individual Web site) enables you to select any of the following logging formats:

- [Microsoft IIS Log File Format](#) – A fixed-format ASCII text log format that provides extensive logging capabilities for Web and FTP sites.
- [NCSA Common Log File Format](#) – A shorter, fixed-format ASCII log format that logs activity to Web, SMTP, and NNTP sites.
- [ODBC Logging](#) – Logs a fixed set of data fields to an Open Database Connectivity (ODBC)-compliant database that can be entered into an Access or SQL Server database.
- [W3C Extended Log File Format](#) – The default log file format for IIS, this is a customizable ASCII text formatted log file that enables you to select the properties to be tracked. To select the properties to be logged, click **Properties** and select the appropriate items in the **Advanced** tab of the log's Properties dialog box.

You can also create custom logging modules that provide more flexible logging options for IIS. See [Custom Logging Modules](#) and references found therein for further information.

The following are several recommended practices when configuring IIS logging:

- Select an appropriate log file format based on your site's requirements. Information on the properties logged by each format is found in the Web links provided in this section.
- Set appropriate file and folder permissions for the logs. Use a GPO linked to the OU in which the IIS servers reside, and configure a policy in the Computer Configuration\Windows Settings\Security Settings\[File System](#) node.
- Configure auditing policies for the logs and for the Web applications, folders, and files. Use the same GPO and configure the appropriate policies in the Computer Configuration\Windows Settings\Security Settings\Local Policies\[Audit Policy](#) node.
- If the Web servers are not part of the domain, configure auditing and file system policies in the Local Security Policy console.
- Ensure that only the required ports are open on the firewall, and configure auditing at the firewall for traffic on the allowed ports.
- Ensure that you have a plan or schedule in place for regular review and archiving of all logs and audit files. If you know the regular patterns of activity at your sites, it will be much easier to spot abnormal activities which may signal incursion attempts.

Designing a Content Management Strategy for Updating IIS Servers

Updating content on a Web site is an ongoing management responsibility for Web masters and other administrators. It is essential to have a strategy in place that enables only those individuals who are responsible for updating content to have that ability. This is an important component of keeping your IIS servers secure.

The method you enable for updating content depends on conditions such as Web site applications, firewall configurations, methods used to create Web site content, and so on. You can choose from the following methods for updating your content:

- [FTP](#) – This is a simple method for uploading content to a Web site; however, it must be secured very carefully to prevent unauthorized individuals from modifying content. You should ensure that TCP ports 20 and 21 are closed on the external firewall; open these ports on an internal firewall only when required for uploading content. Ensure that anonymous FTP connections are not permitted, and that appropriate FTP site and NTFS permissions are present on all folders and files. If a Web server hosts more than one Web site and different individuals are responsible for updating content, consider using [FTP user isolation](#) to control the update process.
- [WebDAV](#) – This is an extension to the basic HTTP protocol that uses the same authentication methods already present on the Web server. You can use SSL for encrypting data uploaded to the site.
- [Microsoft FrontPage Server Extensions](#) – If you are using FrontPage and related Microsoft products as a Web site development tool, you can update content in a method similar to that used with WebDAV. You can also use SSL for data encryption.
- [Use of file shares](#) – This is a simple means of updating content that relies entirely upon shared folder and NTFS permissions for its security. It uses the Server Message Block (SMB) protocol on TCP ports 138 and 139. For a Web server in a perimeter network, open these ports only when required for uploading content. In addition, use IPSec policies to authenticate clients and secure data being sent to the Web server.

Designing Security for Communication Between Networks

Almost every company does business from different offices in different locations. Consequently, networks at these locations must communicate with each other across the Internet or other connections that may not be secure. You can ensure that communications between networks remain secure by establishing a VPN for connection between the networks.

Selecting Protocols for VPN Access

Windows Server 2003 provides the following protocols for VPN access:

- [Point-to-Point Tunneling Protocol \(PPTP\)](#) – Uses Microsoft Point-to-Point Encryption (MPPE) to encrypt Point-to-Point Protocol (PPP) traffic being sent across a tunneled IP network connection. Uses the RSA/RC4 algorithm with 40-, 56-, or 128-bit encryption keys to encrypt only the data portion of the packet.
- [Layer Two Tunneling Protocol over IPSec \(L2TP/IPSec\)](#) – Encapsulates packets including UDP, L2TP, and PPP headers, and then wraps these with an IPSec ESP header and trailer that provides mutual authentication, data confidentiality, integrity, and nonrepudiation. This protocol uses DES or 3DES for encryption and requires computer certificates for the computers at each end of the tunneled connection.

- [IPSec tunnel mode](#) – Creates a secure encrypted tunnel between routers or gateways that do not support either PPTP or L2TP/IPSec. You can use this protocol to connect to a network that uses a third-party IPSec gateway, an IPSec-enabled router situated in front of a third-party server, or across a third-party firewall.

Keep in mind the following considerations when selecting a protocol for securing data sent across a VPN:

- L2TP/IPSec requires computer certificates for each end of the tunnel, whereas PPTP does not. To use L2TP/IPSec, you must either have a PKI established or purchase computer certificates from a third-party agency. If neither is practical, you should use PPTP.
- You cannot use NAT with L2TP/IPSec unless the VPN client and server computers are configured for NAT Traversal (NAT-T), which is a new feature of NAT in Windows Server 2003. When used with IPSec encryption, Microsoft VPN servers prior to Windows Server 2003 think that packets crossing the NAT device have been improperly modified; therefore, these packets are dropped. To allow use of NAT-T with computers running Windows NT 4.0, Windows 98, or Windows 2000, install the [Microsoft L2TP/IPSec VPN client](#).
- Of the protocols discussed, L2TP/IPSec provides the greatest level of security because it provides strong end-to-end encryption as well as mutual authentication of users and computers at both tunnel end points.
- If you are having problems getting L2TP/IPSec to work, try using PPTP to test the VPN connectivity. If the connection works using PPTP, the problem is related to protocol configuration; if it doesn't work, the problem is most likely in the connection.
- You should use IPSec tunnel mode only if one of the tunnel end points does not support L2TP/IPSec. IPSec tunnel mode does not provide the additional security of mutual authentication of users and computers at each end.

Designing VPN Connectivity

When designing a VPN, consider the location of the VPN server on the network as well as the authentication, authorization, and encryption to be used. Refer to [Enterprise Design for Remote Access](#) for details of VPN design.

Position the VPN server in any of the following locations:

- *In front of the firewall* – Enables you to configure the server to accept only traffic related to the VPN protocol in use. The VPN server decrypts traffic and passes it to the firewall, and you can configure the firewall to filter traffic according to the resources that should be allowed.
- *Behind the firewall* – Enables you to configure the firewall to allow the VPN traffic (PPTP and/or L2TP/IPSec) and direct it to the VPN server. The firewall provides filtering and logging of VPN access traffic, and provides tight security for the VPN server.
- *Beside the firewall on the same network segment* – Provides additional routing functionality but leaves the VPN server more vulnerable to attack.
- *Consolidated with the firewall* – A cost-effective, manageable solution but may cause conflicts in delivery of firewall and VPN services.

Table 1 provides information on available user [VPN authentication protocols](#):

| Protocol | Description | Usage |
|--|---|--|
| Password Authentication Protocol (PAP) | Sends passwords in plaintext format. | Do not use unless no other protocol is suitable. |
| Challenge Handshake Authentication Protocol (CHAP) | Sends an MD5 hash of the password. Server needs to store the password using reversible encryption. | Use only when UNIX clients are present. |
| Microsoft CHAP (MS-CHAP) | Uses an MD4 hash and enables the server to store a hashed password. | Use only if Windows 95 clients are present. |
| MS-CHAP v2 | Improves on MS-CHAP by providing mutual authentication of server and client. | Can use with Protected EAP (PEAP) for authenticating wireless clients. |
| EAP-Transport Layer Security (EAP-TLS) | Enables mutual authentication based on a public key certificate that can be stored on a smart card. | Required if using smart cards for authentication. |
| EAP-MD5 | Similar to EAP-TLS but does not use certificates and does not generate cryptographic keys. | Use with reversibly encrypted passwords on L2TP/IPSec VPNs and dial-up access. |
| EAP-MS-CHAP v2 | Password-based and provides mutual authentication. Both client and server must know passwords. | Use with PEAP. Does not require a PKI. |

Table 1 – VPN Authentication Protocols

Designing Demand-Dial Routing Between Internal Networks

Demand-dial routing refers to the initiation of a connection between two networks when data needs to be transferred between the networks. The server will dial out to an ISP or to the phone number of a modem in the remote office. You can set up demand-dial routing by configuring Routing and Remote Access (RRAS) on the servers at each end of the connection. For a review of demand-dial routing, refer to [What Is Demand Dial Routing?](#) and [How Demand Dial Routing Works](#).

When designing a secure demand-dial connection between two offices, observe the following suggestions (see also [Demand-dial routing design considerations](#) for more information):

- On each RRAS server, create a user account whose name matches the name of the interface at the other end of the demand-dial connection. The remote end of the connection uses this account to authenticate to your RRAS server. Clear the **User Must Change Password at Next Logon** option and select the **Password Never Expires** option. Choose a long, complex password for this account. You should also change the password regularly and coordinate this password change with the other office.
- For a two-way, demand-dial connection (either end of the link can initiate a connection), repeat these configuration actions at both ends.

- Use the strongest means of authentication and encryption that is compatible with your network components. Although you could use any of the authentication protocols discussed in the previous section, use EAP-TLS if a PKI is available; otherwise, use MS-CHAP v2.
- Ensure that certificate revocation checking is enabled and that the latest certificate revocation list (CRL) is always available.
- In an installation that uses multiple VPN servers (including demand-dial and RAS servers), configure a Microsoft IAS server for centralized authentication and accounting.
- Use L2TP/IPSec encryption for strong end-to-end encryption and strong authentication. Also configure the firewall to allow L2TP/IPSec traffic to pass.
- Configure packet filtering on each interface of the connection to restrict the traffic that is allowed across the link.
- For added control over the use of the demand-dial connection, configure remote access policies that specify considerations such as the types of connections allowed, the Windows groups allowed to make a connection, and the hours during which a connection is permitted.
- Even if you are using a leased line, such as an ISDN line, rather than the Internet to make the connection, you should still follow the suggestions presented here. Eavesdropping and man-in-the-middle attacks can still occur over the leased line.

Designing Security for Wireless Networks

Wireless networks are becoming more important as they provide enhanced mobility and ease of connection for users who need to access networks and the Internet from diverse locations.

Operating according to the Institute of Electrical and Electronics Engineers (IEEE) 802.11 series of standards, wireless networks can operate in either of [two modes](#):

- *Ad hoc* – Refers to a direct wireless connection between two computers equipped with wireless network interface cards (NICs)
- *Infrastructure* – Refers to a wireless network containing an access point connected to a regular LAN, including products tested and certified by the [Wi-Fi Alliance](#)

Designing Public and Private Wireless LANs

Wireless LANs (WLANs) provide a simple means of setting up network connectivity, particularly in spaces where client mobility is an important factor. See [How 802.11 Wireless Works](#) for an overview of 802.11 technology. When designing the WLAN, you can choose between the following authentication methods:

- *Open system authentication* – Uses device-oriented authentication that enables connection of a wireless device to an access point, according to knowledge of the correct service set identifier (SSID) and wired equivalent privacy (WEP) key.
- *Shared key authentication* – A wireless device can connect to the network by providing an alphanumeric character string known as a shared key or shared secret. This method is vulnerable to interception of the key by an intruder.

When designing security for a WLAN, you should take the following into consideration (See [Wireless Networking Security](#) for more details and additional recommendations):

- Attackers can break into your network if they can sniff the SSID and WEP keys. You need to change the default SSID and disable its broadcast on the network. You also need to plan a secure method for distributing WEP keys.
- If using WEP, configure it for the highest possible security. Use 128-bit keys rather than 40-bit ones, and ensure that you have configured WEP for dynamic rekeying.
- Use more than WEP for encrypting your data. You can (and should) include IPSec, VPN, and SSH in your WLAN design wherever possible.
- Use [Wi-Fi Protected Access](#) (WPA) to improve WLAN security by requiring 802.1x authentication, encryption, and data integrity. We discuss 802.1x in the next section.
- Attackers often scout for available WLANs by war driving—in other words, driving around with a laptop and wireless card looking for available networks. Design your access points to minimize the potential from access outside your premises, and perform a [site survey](#) to check the extent to which your WLAN is available.
- Employees or visitors might attach a rogue access point to your network. Your site survey should be capable of checking for the existence of such points and include periodic scans to ensure new ones haven't been installed since your initial survey.
- Many access points use [Simple Network Management Protocol](#) (SNMP) for configuration and monitoring purposes. You should change the community string on the access point so that an intruder cannot easily guess it.
- Use a firewall or packet filter to separate the WLAN access points from the remainder of the network.

Designing 802.1x Authentication for Wireless Networks

The IETF developed 802.1x to act as a solution for the security problems inherent with the 802.11 wireless standards. 802.1x provides a means of port-based network access control as it applies to both wired and wireless communications. It works Microsoft IAS to provide EAP-MS-CHAP v2 or EAP-TLS as already discussed in Table 1 for mutual authentication of clients and servers on the WLAN. (See [Designing the Wireless LAN Security Using 802.1x](#) for more information). 802.1x can also provide EAP-PEAP, which provides a dedicated encryption channel, dynamic keying material generated from TLS, fast reconnection when roaming between access points, and server authentication that can protect against the use of unauthorized wireless access points. You can also configure media access control (MAC) filtering on many access points to limit access to only approved wireless clients.

Group Policy in Windows Server 2003 provides the Wireless Network (IEEE 802.11) Policies policy located under Computer Configuration\Windows Settings\Security Settings. Right-click this node and choose **Create Wireless Network Policy** to create a policy and configure its properties. Use the Preferred Networks tab of the policy's Properties dialog box to specify the required network parameters including the variables (*Figure 15*) that define the 802.1x authentication and authorization parameters. See [Define 802.1x authentication for wireless networks in Group Policy](#) for information on configuring the available options.

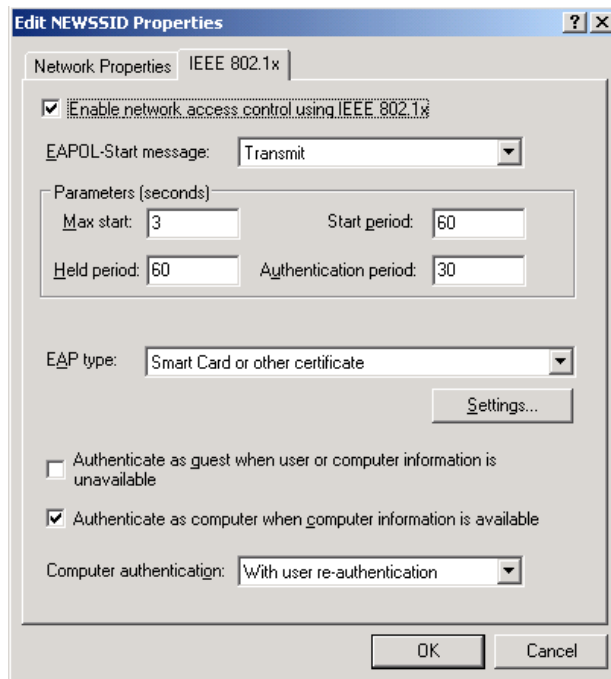


Figure 15 – The IEEE 802.1x Tab of the WLAN's Properties Dialog Box Enables You to Configure Access Parameters

The following are several guidelines you should follow when designing your wireless network policy. See [Design Considerations for Wireless Network Policies](#) and [Designing On a Secure Wireless Networking Strategy](#) for additional tips when designing your wireless policy:

- Configure wireless policies in GPOs linked to various OUs to specify conditions for connecting to the network for different departments or work groups as required.
- Use an IAS server to authenticate your users, and configure the access points as IAS clients. This enables you to define the appropriate remote access policies for connecting to the WLAN. You can also restrict access to the access points and control parameters such as encryption strength and IP packet filters.
- EAP-TLS provides the most secure authentication method on networks containing a PKI. Ensure that you use CRL checking so that intruders cannot use stolen laptops containing access credentials for attacking your network.
- If your network does not use PKI, use PEAP-EAP or EAP-MS-CHAP v2 for authentication.
- When remotely administering the wireless access points, use SSL or SSH to reduce the usefulness of intercepted data transmissions.

Designing Security for Communication with External Organizations

Designing an Extranet Infrastructure

Modern business needs include the confidential exchange of information between related organizations such as companies that have formed partnerships, client companies, trusted customers, and so on. Common methods of communication with external organizations include VPNs and Web applications secured with SSL.

- You can set up an *extranet*, which is a Web-based connection that operates only between the companies involved. Client certificates from each company are mapped to the appropriate Active Directory accounts, and are used to authenticate requests to the Web application. SSL is used to secure communications between Web applications operated by the partner companies, and IPsec is used to secure internal database communications. See [Securing .NET Web Applications in an Extranet Environment](#) for details in designing a Web-based extranet setup.
- You can use certificate-secured communications across a VPN. Certificates used in these communications must be trusted by all external organizations. To accomplish this, the root certificates from the CA hierarchies issuing the certificates must be present in the validating computer's Trusted Root Certification Authorities certificate store. Alternately, you can purchase all certificates required from a common public CA such as VeriSign. This enables all certificates to automatically be trusted by the various organizations. See [Virtual Private Networking with Windows Server 2003: Deploying Site-to-Site VPNs](#) for details in designing a VPN-based communication with an external company.
- You can also set up a cross-certification infrastructure between the CAs that are managed by different organizations. We discuss cross-certification of CAs next.

Designing a Strategy for Cross-Certification of Certificate Services

We discussed designing a PKI implementation using Certificate Services earlier in this exam manual. [Cross-certification](#), also known as *qualified subordination*, refers to the enabling of trust between the certificates of each of two or more organizations, each of which operates a separate CA hierarchy. It uses a process called *certificate chaining*, which defines the path from any given certificate back to the root CA in the hierarchy that originally issued the certificate.

You create a cross-certification by issuing a certificate based on the Cross Certification Authority certificate template from a CA running Windows Server 2003. This certificate is issued from the root CA of one hierarchy to the root CA of the other; to provide complete cross-certification of both hierarchies, a certificate must be deployed in both directions. You can also issue the certificate between intermediate CAs when deeper CA hierarchies exist. This provides the following benefits:

- You can make certificate revocation information available more rapidly, because subordinate CAs often issue CRLs more frequently than root CAs.
- You can set up a limited cross-certification trust relationship between portions of two CA hierarchies.

Another means of cross-certifying multiple CA hierarchies is to deploy a *bridge CA*, which acts as a link between the hierarchies involved. The bridge CA is cross-certified to all participating CA hierarchies, and simplifies the process of cross-certification when three or more hierarchies are involved. A common scenario involving the use of a bridge CA is a company with several independent subsidiaries.

Designing Security for Servers with Specific Roles

Servers performing different roles require different security settings. Web servers running IIS or e-mail servers running Microsoft Exchange Server 2000 or 2003 that are exposed to the Internet have different security requirements from internal servers, such as domain controllers, infrastructure servers (DHCP, DNS, or WINS), file servers, IIS servers, terminal servers, and mail servers.

Microsoft provides the [Windows Server 2003 Security Guide](#), which is a comprehensive guide in PDF format that guides you through creating security for all types of Windows Server 2003 computers according to the security environment representative of your organization's needs. This guide defines the following three environments, all of which are applicable to an Active Directory-based network:

- *Legacy Client* – Includes member servers and client computers running Windows 98, Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. Domain controllers must be running Windows 2000 or higher.
- *Enterprise Client* – All servers and client computers run Windows 2000 or higher, and the network operates at a medium security level.
- *High Security* – Similar to Enterprise Client, but imposes a higher level of security on all systems. Settings are so restrictive that some applications may not function, or performance may suffer. These settings require a higher level of testing and verification than those in the Enterprise Client environment.

Defining a Baseline Security Template for All Systems

A [security template](#) is an .inf file containing a series of settings that you can apply to one or more computers to define the settings contained within the Computer Configuration\Windows Settings\Security Settings node of Group Policy. You can use any of several methods to apply security templates:

- *Group Policy* – By applying the template to a GPO, its settings automatically apply to all computers subject to the GPO. Right-click **Security Settings** and choose **Import Policy**, and then select the required template.
- *Security Configuration and Analysis* – This MMC snap-in enables you to apply the settings defined in a template to one or more computers. Right-click **Security Configuration and Analysis** and choose **Configure Computer Now** to configure the computer with the settings contained in the specified template.
- *The Secedit command* – By using **secedit** with the **/configure** parameter, you can apply settings in the specified template to one or more computers. You can use a script to apply these settings; this is especially useful for configuring stand-alone computers that are not members of a domain. Refer to [Securing Stand-Alone Windows XP Clients](#) for more information.

Windows Server 2003 provides a series of predefined security templates which you can use as-is or customize according to the needs of your network. You can copy any of these templates and then use the Security Templates snap-in to apply the required customized settings. See [Design Recommendations for Using Predefined Security Templates](#) for suggestions on the use of these templates.

Neither the Security Templates nor the Security Configuration and Analysis snap-ins are included within a default console. It is suggested that you open a blank MMC and add these templates to the console. Doing so provides you with a console (*Figure 16*) that you can use for configuring and administering security templates.

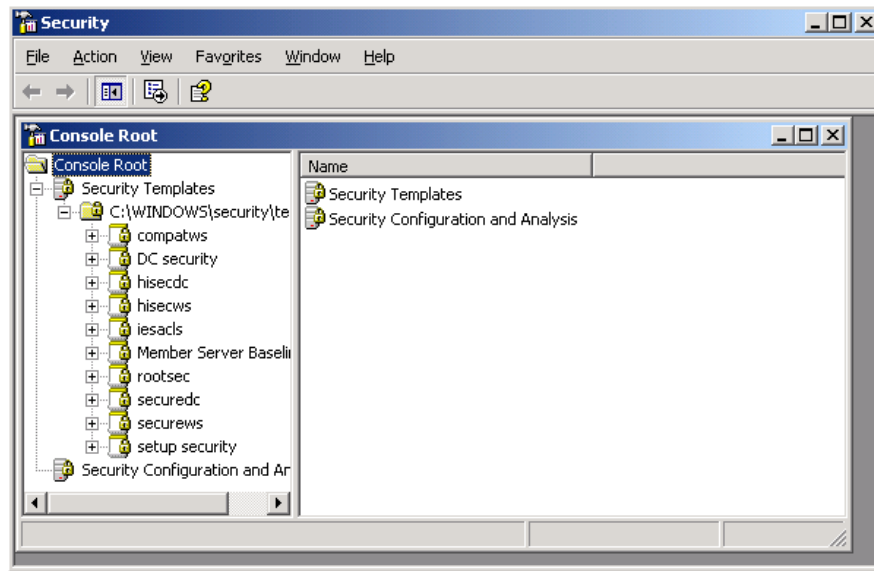


Figure 16 – Add the Security Templates and Security Configuration and Analysis Snap-Ins to a Blank MMC to Create a Console for Administering Security in Windows Server 2003

Using Group Policy to apply security templates to servers according to role requires the existence of the appropriate Active Directory structure. By default, all domain controllers are located in the Domain Controllers OU, and all member servers are located in the domain root. You should create an OU structure that includes a Servers OU for all member servers and a child OU for each server role beneath this OU, and then place the servers in the appropriate OUs. See *Figure 17* for an example. This structure enables you to create one GPO that applies blanket settings that are applicable to all member servers (linked to the Servers OU) and additional GPOs that apply role-specific settings to the OUs containing servers with the indicated role.

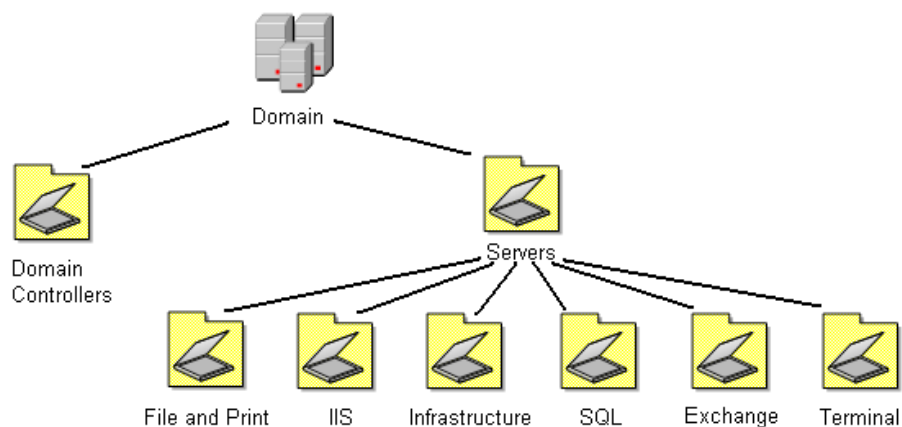


Figure 17 – Sample OU Structure for Domain Servers


Chapter 3 of the Windows Server 2003 Security Guide referenced earlier discusses the settings contained in a Member Server Baseline Policy, which should be linked to the Servers OU in the example shown in *Figure 17*. Included with the Security Manual are sample baseline templates for member servers in the Legacy Client, Enterprise Client, and High Security environments. You can copy one of these baseline templates to the %systemroot%\Security\Templates folder and then modify any settings required to customize the template according to the needs of your company.

Creating a Plan to Modify Baseline Security Templates According to Role

The Windows Server 2003 Security Guide contains sample templates that apply role-specific settings to servers in the Legacy Client, Enterprise Client, and High Security environments. Included are templates for the file server, print server, infrastructure server, IAS server, certificate server, bastion host, and IIS server roles. As with the member server baseline templates, you should customize these templates to fit the security needs of your organization, and create additional templates for roles, such as Terminal Server, that are not included in the Security Guide. You would then import each template to the Security Settings subnode of a GPO linked to the appropriate child OU.

The following are some of the modifications that you should consider when customizing security templates:

- *Domain controllers* – The Security Guide contains domain controller templates for the three security environments. These templates contain settings for improving the security for user rights as they affect access to domain controllers, system services related to DNS, file replication (FRS), Distributed File System (Dfs), intersite messaging, Kerberos, and Remote Procedure Call Locator, and several other security settings (see Chapter 4 of the Security Guide for details). A suggested set of IPSec filters is also included. You should customize this template according to your organization's needs and import it to the Default Domain Controllers Policy GPO or to another GPO linked to the Domain Controllers OU.
- *Infrastructure servers* – The incremental template included with the Security Guide enables the DHCP and WINS services, protects against DoS attacks, and secures well-known accounts. It also includes a suggested set of IPSec filters.
- *File servers* – The incremental template disables the Dfs and FRS services, secures well-known accounts, and provides a set of IPSec filters. If you are using either of these services, you must modify the template. You may also need to configure file system permissions as required by your setup.
- *IIS servers* – The incremental template restricts internal access to IIS servers, and enables services related to IIS. Chapter 8 of the Security Guide also provides guidelines for IIS subcomponents you should install under various conditions, NTFS permissions, IIS permissions, and IPSec filters.
- *Terminal servers* – Not included with the Security Guide, you should configure the Terminal Services service for automatic startup, and configure mirrored IPSec filters that allow traffic on the proper ports and protocols required for user terminal sessions. You should also configure a [Restricted Groups policy](#) for the Remote Desktop Users group. If the terminal servers operate as a cluster, you should also configure the Terminal Services Session Directory service for automatic startup.
- *POP3 mail servers* – Also not included with the Security Guide, you should configure the POP3SVC service for automatic startup, and configure mirrored IPSec filters for allowing traffic based on POP3 and ports and protocols required for communication with domain controllers and blocking traffic from nonessential protocols and services. Refer to [Microsoft Exchange Server 2003 Hardening Guide](#) for recommendations on additional service startup types.



Ready to pass the 70-298 exam?

Download a **free** practice exam preview to find out if you're ready to pass.

Designing an Access Control Strategy for Data

One of the major defenses against an attack on your network, either internal or external, is the access control mechanism provided for all files, folders, and directory objects. Permissions at the object level enable you to control access at the granular level and protect resources from access or modification by unauthorized individuals, including employees. In this section, we look at methods of designing access control strategies and then review methods of auditing data access.

Designing an Access Control Strategy for Files and Folders

One of the major cornerstones of security in any network is granting proper access control to resources, such as files and folders, on all computers. Even if an intruder manages to access a computer by cracking a user account, he will be unable to access any resources to which the user has not been granted permission.

All permissions for various objects, including files, folders, registry keys, Active Directory objects, and so on, are defined in access control lists (ACLs), which in turn contain access control entries (ACEs).

Each security descriptor contains the following types of ACLs:

- *Discretionary ACLs (DACLS)* – Identify the users and groups that have been specifically granted or denied access to the object.
- *System ACLs (SACLs)* – Identify the users and groups whose access to the object is to be tracked. We discuss auditing requirements later in this exam manual.

Designing a Permission Structure for Files and Folders

When designing a permission structure for files and folders in Windows Server 2003 you must be aware of the default permissions provided by Windows Server 2003, how these permissions differ from those in previous Windows versions, and how permissions interact. If you need to refresh your knowledge of Windows Server 2003 access controls, refer to the links provided at [Access Control Concepts](#). In particular, keep the following facts in mind:

- The default shared folder permission in Windows Server 2003 has been changed to Everyone, Read. Remember that shared folder permissions affect access only when the resource is accessed across the network, and that users receive the most restrictive set of permissions between the shared folder and NTFS permissions.
- The default NTFS permission in Windows Server 2003 has been changed to Users, Read & Execute, List Folder Contents, and Read. In addition, Administrators, the SYSTEM group, and the Creator Owner of an object receive full control. Remember that all users receive the cumulative extent of all permissions provided them from group memberships unless explicitly denied access.

Keep in mind the following when designing an access control strategy for files and folders (see [Additional Resources for Designing and Deploying File Servers](#) for additional information):

- Always ensure that all volumes are formatted with the NTFS file system. Otherwise, you will be unable to assign file and folder permissions.
- By default, permissions set at the topmost level of a folder hierarchy are propagated throughout the hierarchy. You can change this behavior by clearing the **Allow inheritable permissions from the parent to propagate to this object and all child objects** check box, found on the Advanced Security settings dialog box for the parent folder. When designing security for sensi-

tive material, it is a good practice to clear this option so that such material does not inherit the parent folder (or volume) permissions.

- If you have a large number of folders and files that require the same permissions, organize these into a single folder hierarchy and apply the required permissions to the root folder.
- Assign permissions to groups rather than to individual users.
- Try to avoid using explicit denial of permissions. Remember that such a denial overrides all allowed permissions, and can prevent users from accessing material required in the performance of their job tasks.
- Do not modify permissions on system files and folders unless absolutely necessary. These are set by default to provide the best security without loss of performance.
- If users are having problems accessing resources, remember that the Effective Permissions tab of the Advanced Security Settings dialog box for any file or folder summarizes permissions granted to the user or group in question.

Designing an Encryption and Decryption Strategy

Encrypting File System (EFS) provides an additional layer of security for files and folders in Windows 2000, Windows XP, and Windows Server 2003. It is especially valuable in computers such as laptops that leave the premises and are at risk for theft. See [Data Protection and Recovery in Windows XP](#) for an overview of EFS and how it works to provide data security.

Keep in mind the following when designing an encryption and decryption strategy:

- To ensure that sensitive data is encrypted as it is created, configure the folders in which this data is stored for encryption, rather than the individual files.
- Ensure that temporary folders used by applications are also encrypted. Otherwise, these folders could contain unencrypted copies of files.
- Remember that EFS does not encrypt data transmitted across the network. Use IPSec for securing data during transmission.
- You can also use WebDAV to store and transfer EFS-encrypted files and folders on servers running Internet Information Services (IIS). For more information, see [Planning Encrypted File Storage](#).
- You should design a policy in Group Policy for use of EFS. Navigate to Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Encrypting File System. From this location, you can add and create data recovery agents, or disable the use of EFS entirely.
- Remember that you need to archive and store recovery keys used by EFS in a secure manner. Keep in mind the following facts about data recovery:
 - ▶ You should export data recovery keys from the computer on which EFS is used, and store them securely. If an intruder has possession of the recovery keys, she can decrypt all files they were supposed to protect.
 - ▶ Back up all EFS recovery keys and store them in a secure place such as a vault.
 - ▶ Set up a data recovery station on a dedicated computer. You can protect access to this computer and use it only for data recovery. When required, files can be copied to this computer, decrypted, and then returned to their owner.

- ▶ If no data recovery agent is available and a user loses his EFS decryption keys, all files and folders encrypted by this user become permanently unavailable.
- ▶ Create and distribute a policy for encrypted file recovery. You should include the guidelines and steps for file recovery and the standard operating procedures for storing and obtaining recovery keys. Ensure that users are educated in the use of EFS and its proper usage. This should include the creation of a password reset disk when changing passwords, in order to avoid permanent data loss should the password be reset by another user (for example, if the user forgets her password).

Designing Security for Backup and Recovery

A secure strategy for backing up and restoring corporate data is part and parcel of a company's daily operations. Backup media contains copies of all your company's vital data and a thief could use your backup media to restore to her own servers and thereby steal all your vital information. You should observe the following guidelines to ensure that your backup and restore operations are properly secured (see [Backup and Recovery Services Design](#) and [Best Practices for Backup](#) for more information):

- Use secure communications strategy when backing up remote servers including packet filtering and port blocking. In addition, backup software should log all failed access attempts. We discuss auditing and logging later in this section.
- Ensure that backup media are properly labeled and protected. Strategies include password protection, media rotation, and storage both onsite and offsite. Secure storage locations include fireproof vaults, locked safes, and bank safety deposit boxes. Companies also exist that provide offsite storage of backup media in controlled vaults and the like. They can also provide services such as restoration of data offsite and deliver the data on a different medium that's easier to manage.
- Implement an access control procedure for backup activities. Separate the backup and restore privileges (by default provided to the Backup Operators group) so that different individuals are empowered to perform backup and restore operations.
- Companies often use a storage area network (SAN) as a storage location for data backups. You should ensure that transfer of data and its storage in the SAN are conducted in a secure manner.
- All backup media has a finite lifetime. You should periodically perform test restores to ensure that the backup procedure and media are performing properly. You should ensure that worn out media is destroyed in a secure manner. Refer to Sun Microsystem's [Data Security Policy - Structure and Guidelines](#) document for information on secure destruction of backup and other media.

For a case study-based review of a secure backup design strategy, refer to [Storage and Backup Services](#). For an example of secure backup involving Exchange servers, refer to [Exchange 2003 Design and Architecture at Microsoft](#).

Designing an Access Control Strategy for Directory Services

As already mentioned, each object in Active Directory has an ACL associated with it. To view and modify these ACLs, open Active Directory Users and Computers, and then click **View > Advanced Features**. Right-click any object and choose **Properties**. Select the **Security** tab to view and edit permissions, as shown for an OU named Accounting in *Figure 18*. For more information on ACLs and a general introduction to access control in Active Directory, refer to [Access Control in Active Directory](#).

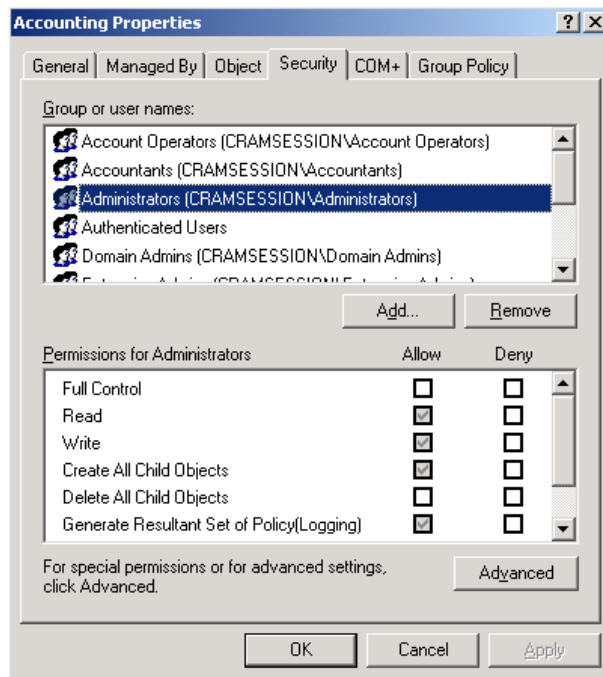


Figure 18 – The Security Tab of an Object's Properties Dialog Box in Active Directory Enables You to View and Modify Permissions

Creating a Delegation Strategy

Active Directory enables you to delegate control of a portion of the Active Directory namespace to specific users or groups. The [Delegation of Control Wizard](#) enables you to assign permissions for performing only those tasks required for users to fulfill their work duties, thereby enabling you to enforce the principle of least privilege, which means that a user should receive only the minimum amount of user rights and permissions required to perform his job tasks.

As an example, the following steps show you how to delegate the task of resetting passwords in the Accounting OU to a group named Support:

1. In Active Directory Users and Computers, right-click the Accounting OU and select **Delegate Control**. This starts the Delegation of Control Wizard.
2. Click **Next**, click **Add**, and type **Support** in the Select Users, Computers, or Groups dialog box.

- Click **OK** and then click **Next** again.
3. On the Tasks to Delegate page (Figure 19), select **Reset user passwords and force password change at next logon**, and then click **Next**.
 4. On the Completing the Delegation of Control Wizard page, review the information provided and then click **Finish**.

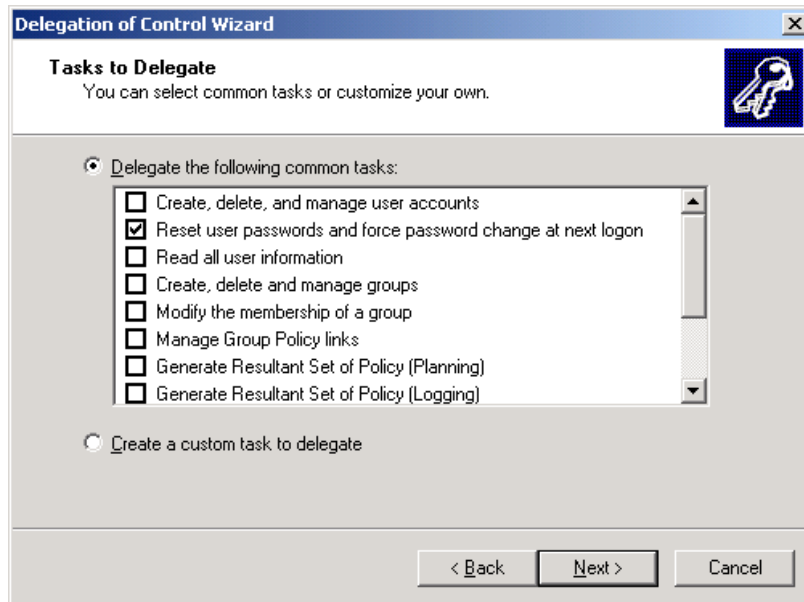


Figure 19 – Delegating the Reset Passwords and Force Password Change Tasks

The following are several best practices you should follow when delegating administrative tasks (for additional details see [Best Practices for Delegating Active Directory Administration](#)):

- Design the Active Directory group and OU structures with the requirements for delegating authority in mind.
- Do not assign the Full Control permission unless absolutely necessary.
- You cannot use the wizard to remove control from a group. If you need to remove a delegated permission, open the Security tab of the OU's Properties dialog box and remove the permission assigned by the wizard.
- Delegate authority to groups rather than to users. This makes it simple to add or remove authority by adding or removing users from the respective groups.
- Create a custom console that includes only the tasks that you have delegated. You can create a [taskpad view](#) that simplifies usage of the console by the group to whom you have delegated the task in question.
- When designing an administrative control scheme, test your proposed design in a lab environment before deploying it to the production network.

Designing a Group Strategy for Accessing Resources

Windows Server 2003 provides extensive capabilities for creating and managing groups. You can create the following three security [group scopes](#):

- *Universal groups* – These groups exist across multiple domains within a single forest. They can contain users or groups from anywhere in the forest, and can be assigned permissions to objects in any domain in the forest. They are available only if the domain functional level is set to Windows 2000 native or higher.
- *Global groups* – These groups exist within a domain and can contain users and groups from the domain in which they are created. You can assign permissions to resources in any domain in the forest.
- *Domain local groups* – These groups exist within a domain and can contain users and groups from any domain in the forest, and can be assigned permissions only to objects within the domain in which they are created.

Windows Server 2003 also provides an extensive list of [default groups](#), which are preconfigured to perform particular types of administrative activities in the domain in which they are created. You can add users directly to these groups where their privileges provide the required administrative capabilities. More frequently, you will need to create new purpose built groups and add users to these groups to design a structure that provides for the required resource access and administrative capabilities.

The following are several best practices you should follow when designing a group strategy:

- Use default groups only when their capabilities match the required privileges.
- Create user and group accounts in their respective OUs, and then delegate control of the OUs to these groups as required.
- Use group nesting to provide the required access. The recommended procedure is to place users in global groups, add these groups to domain local groups in the domain where access is required, and then grant permissions to the domain local groups.
- Do not assign privileges to individual user accounts. Doing so can create a structure that is extremely difficult to manage and troubleshoot.
- Use the principle of least privilege. Assign groups only the minimum amount of access required for users to do their jobs properly.

Designing a Permission Structure for Directory Service Objects

As we have already noted (and shown in *Figure 18*), each object in Active Directory has its own ACL associated with it. When designing a permission structure, you should keep the following [best practices](#) in mind:

- *Avoid changing default permissions* – Doing so could lead to unexpected problems or reduce the default security.
- *Do not grant the Full Control permission if at all possible* – Use the principle of least privilege and assign only the required permissions.

- *Use the fewest possible number of ACEs on child objects* – If you are using the **Apply Onto** option for controlling inheritance, you should be aware that all child objects will receive a copy of that ACE. This could lead to performance problems.
- *Try to assign the same series of permissions to multiple objects* – A feature in Windows Server 2003 called single-instancing enables Active Directory to store only a single copy of an ACL and associate it with multiple objects, thereby improving performance.
- *Try to assign access rights broadly rather than individually* – The fewer the number of ACEs you assign, the better the system performance. See the link provided in this section for further suggestions.
- *Assign permissions to groups rather than users* – This principle holds true for objects in Active Directory as much as for files, folders, and other resources.

Designing an Access Control Strategy for the Registry

The Windows registry contains vital information without which the computer would be unable to start and applications would be unable to run. An intruder could create a DoS condition by modifying the registry to prevent normal operation. In addition, an inexperienced user could cause problems by accessing and modifying portions of the registry in an improper manner. An appropriate access control strategy helps to avoid these problems. Refer to [Securing the registry](#) for registry permission information.

Registry keys and subkeys follow a permissions and inheritance structure similar to that used by files and folders. In brief, you should run **Regedit.exe**, right-click the required key, and choose **Permissions**. This opens a Permissions dialog box similar to that shown for **HKEY_LOCAL_MACHINE\SOFTWARE** in *Figure 20*, from which you can modify permissions as required.

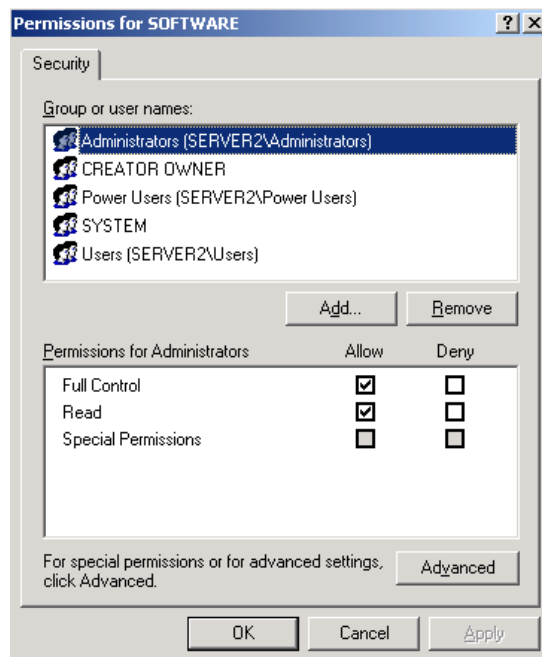


Figure 20 – The Permissions Dialog Box Enables You to Configure Registry Permissions

Group Policy enables you to specify permissions on registry keys for all computers subject to the GPO in which you configure a registry policy. Navigate to the **Computer Configuration\Windows Settings\Security Settings\Registry** node, right-click this node, and choose **Add Key**. In the Select Registry Key dialog box (Figure 21), expand the listing to locate the required entry and then click **OK**. This displays a dialog box similar to that shown in Figure 20, and allows you to configure registry key permissions that apply to all computers subject to the GPO. You can perform the same tasks by editing a security template and applying it to a group policy or to individual computers.

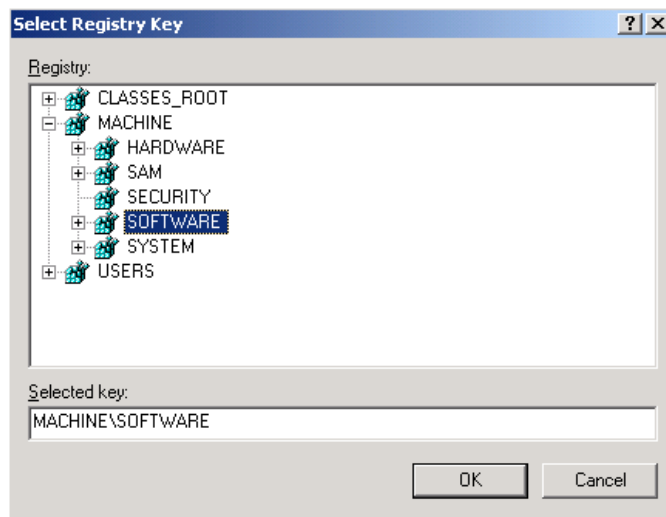


Figure 21 – The Select Registry Key Dialog Box Enables You to Configure Registry Permissions in a GPO

You should not normally need to modify registry permissions. The following are several guidelines you should follow in the event you need to modify registry security:

- Use test computers and networks to ensure that modifications do not disrupt normal computer functionality. This is especially true if you need to grant users additional capabilities for modifying registry keys and values.
- If applications require elevated registry permissions, evaluate and test these thoroughly in conjunction with application developers. Try not to give users administrative rights; often they need permission only to write to the HKEY_LOCAL_MACHINE registry hive.

For further information on registry security, refer to [Microsoft Windows 2000 Security Hardening Guide](#). Although written for Windows 2000, material in this reference is largely applicable to Windows Server 2003 (substitute **Regedit.exe** instead of **Regedt32.exe**).

Analyzing Auditing Requirements for Directory Services, Files and Folders, and the Registry

Having designed access control strategies for directory services, files and folders, and the registry, your next step is to design an auditing strategy. Auditing provides information on attempts by users (whether authorized or not) to access resources on the network, and enables you to spot incursion attempts in progress. Furthermore, legal requirements may dictate a certain level of auditing to ensure that privacy requirements are met.

Windows Server 2003 enables you to define an [auditing policy](#) in Group Policy, either in a GPO linked to an Active Directory container or in the Local Security Settings snap-in. Navigate to the **Computer Configuration\Windows Settings\Security Settings\Audit Policy** node. You can also define an auditing policy in a security template that can be applied using Security Configuration and Analysis or imported into a GPO. All actions that you have enabled for auditing are recorded in the computer's Security log, which you can access from Event Viewer. As shown in *Figure 22*, you can define success or failure auditing for the following actions.

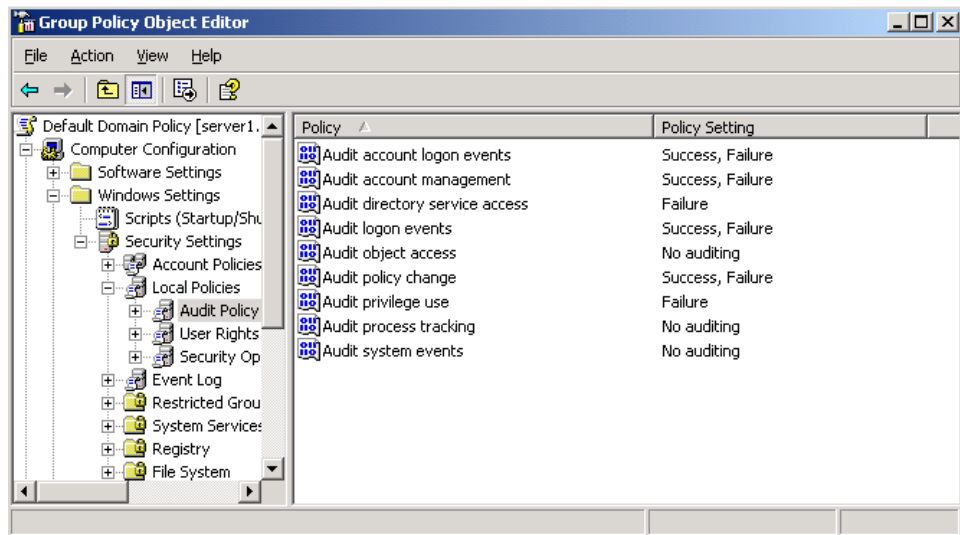


Figure 22 – You Can Configure Auditing for Various System Actions from the Audit Policy Subnode in Group Policy or a Security Template

- *Account logon events* – Attempts at logging onto a computer using a domain user account.
- *Account management*– Creation, modification, or deletion of user or group accounts, or changes in account properties including passwords.
- *Directory service access*– Access of an Active Directory object that contains its own SACL.
- *Logon events* – Attempts at logging onto a computer using a local user account.
- *Object access* – Access of a file, folder, or printer that is configured for auditing.

- *Policy change* – Changes to user rights assignment, audit policies, or trust policies.
- *Privilege use* – A user exercising a user right.
- *Process tracking* – Actions taken by an application such as program activation, process exit, handle duplication, or indirect object access.
- *System events* – Users shutting down or restarting a computer, or events that affect a computer's system or security log.

The following are several factors that you should keep in mind when designing an audit policy:

- Work with managers and others in designing the audit policy. These individuals will have insight into the importance of certain types of information, such as research, legal, or financial data, and the extent of auditing required on such data.
- The amount and type of auditing depends on server role. You can take this into account by configuring auditing on GPOs linked to OUs containing the various types of servers (as shown previously in *Figure 17*).
- If you suspect unauthorized activity is occurring (either by local users or intruders), configure a higher level of auditing for sensitive files or objects, and then remove this auditing at a later time. You should define the types of additional auditing and the circumstances under which to implement these in a written policy before suspicious actions occur.
- It is usually not necessary to audit process tracking except when debugging new applications.
- Remember to have a plan for reviewing security logs on a regular basis.

For a list of recommended best practices to follow when designing and implementing auditing, refer to [Auditing Security Events Best practices](#).

Designing Auditing for Directory Service Objects

Auditing for directory service objects involves configuring success and/or failure for directory service access plus specifying a SACL for each object to be audited.

Windows Server 2003 enables you to audit a large number of directory service objects. For detailed information and suggestions on objects you should audit, refer to [Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations](#).

You use the ADSI Edit snap-in to configure directory service object auditing, as follows:

1. Add the ADSI Edit snap-in to a blank MMC console.
2. In the console tree of the ADSI Edit snap-in, right-click **ADSI Edit** and choose **Connect to**.
3. In the Connection Settings dialog box, select the required naming context (domain, configuration, RootDSE, or schema) and the computer or domain, and then click **OK**.
4. Expand the naming context entry in the console tree to locate and select the object to be audited.
5. Right-click this object and choose **Properties**.
6. On the Security tab of the Properties dialog box, click **Advanced**.
7. On the **Advanced Security Settings** dialog box, select the Auditing tab (see *Figure 23*), and then click **Add**.

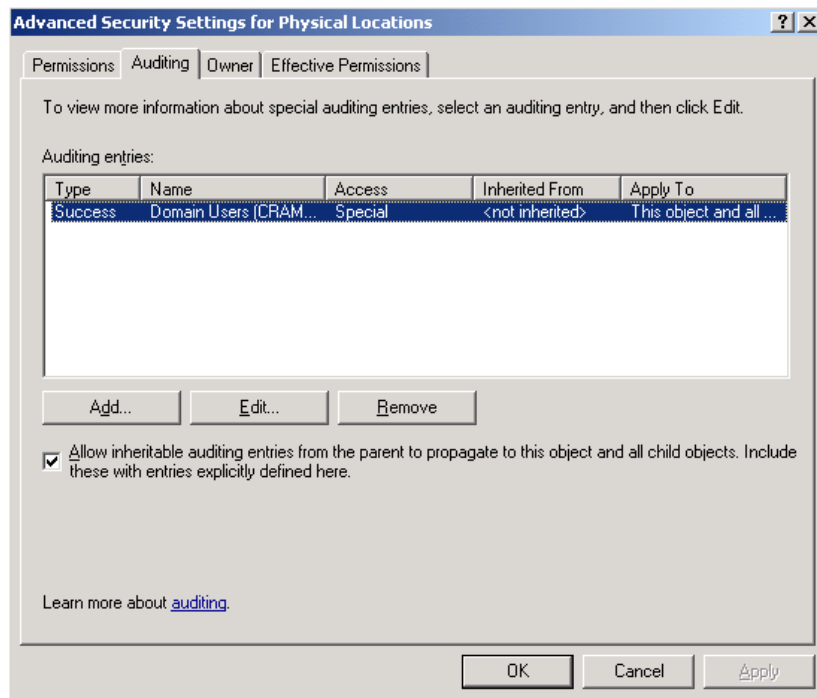


Figure 23 – You Configure Auditing from the Auditing Tab of the Advanced Security Settings Dialog Box for the Object You Want to Audit

8. On the Select User, Computer, or Group dialog box, type the user or group name(s) to be audited, and then click **OK**.
9. The Auditing Entry dialog box opens, enabling you to select the auctions to be audited. Select these under the Successful and/or Failed columns as required, and then click **OK**.
10. Repeat for additional objects as required, and then click **OK** until all dialog boxes are closed.

When configuring auditing for directory service objects (and any other object types for that matter), you should keep in mind the fact that inheritance rules can affect auditing of files, folders, and other objects. You can configure auditing on a parent folder and its subfolders and files will inherit the audit settings. If you don't want subfolders to inherit the audit settings, you can clear the check box labeled **Allow inheritable auditing entries from the parent to propagate to this object and all child objects. Include these with entries explicitly defined here**. This check box is found on the Auditing tab of the Advanced Security Settings dialog box (see *Figure 23*).

Designing Auditing for Files and Folders

Windows Server 2003 enables you to configure [operations-based file and folder auditing](#). In other words, you can audit specific actions such as writing or execution on files and folders. For a general introduction, refer to [Auditing settings on objects](#).

To configure auditing for files and folders, configure success and/or failure auditing for [object access](#) first, as shown previously in *Figure 22*. Then, configure an SAFL for the files and folders to be audited. You can do this by accessing the Security tab of the object's Properties dialog box. Then follow a procedure similar to steps 6 to 10 of that outlined for auditing of directory service objects. For additional details, see [Apply or modify auditing policy settings for a local file or folder](#).

You can also use Group Policy to configure auditing for files and folders. To do so, perform the following steps:

1. In the Group Policy Object Editor, navigate to the **Computer Configuration\Windows Settings\Security Settings\File System** node.
2. Right-click this node and select **Add File**.
3. In the **Add a file or folder** dialog box, navigate to the appropriate folder and then click **OK**. The object's Properties dialog box opens.
4. On the Security tab, configure success and/or failure auditing for [object access](#).

Configuring auditing in this fashion enables you to configure auditing for all computers subject to the GPO in which you perform this configuration.

Designing Auditing for the Registry

Configuring auditing for registry keys is similar to files and folders. First, enable auditing of object access, as already described. Run **Regedit.exe**, right-click the key to be audited, and choose **Permissions**. On the Permissions dialog box (shown previously in *Figure 20*), click **Advanced**, and then select the Auditing tab. Click **Add** to open a dialog box similar to that shown in *Figure 23*, specify the name of the user or group whose actions are to be audited, and then select the actions to be audited from the Auditing Entry dialog box that appears. For a description of the auditable actions, refer to [Audit activity on a registry key](#).

Also similar to the case for files and folders, you can use Group Policy to configure auditing for the registry on all computers subject to the GPO in which you perform this configuration. Access the **Computer Configuration\Windows Settings\Security Settings\Registry** node. Right-click this node and choose **Add Key**, and then follow the same procedure used with **Regedit.exe**.

Creating the Physical Design for Client Infrastructure Security

All computer networks contain large numbers of client computers that must be secured against attack. Besides traditional desktop computers, various portable clients such as laptop computers, personal digital assistants (PDAs), cellular phones, Blackberry devices, and so on, connect to modern networks and need to be secured. In this section we look at securing and authenticating these devices.

Designing a Client Authentication Strategy

A client authentication strategy needs to take into account this diverse set of clients that will be authenticating to your network. In addition, many networks have diverse types of servers such as various flavors of UNIX or even Macintosh Servers, to which clients must be authenticated. You need to identify all such devices and the need for various types of connection points such as remote access and wireless. Furthermore, you need to integrate your authentication design with corporate security requirements as stipulated by the company's written security policy. For a case study-based review of designing client authentication, refer to [Platform and Infrastructure](#).

Analyzing Authentication Requirements

When analyzing requirements for client authentication to your network, you need to keep in mind the various types of clients that might need authentication, including Windows and non-Windows clients located on the LAN, clients accessing the network remotely by means of a VPN or through a secured web-based connection, wireless clients, and so on. We have introduced these various means of authentication earlier in this exam manual. The following are several factors you need to take into consideration when analyzing client authentication needs (see [Network access authentication and certificates](#); refer to [Deploying PKI Inside Microsoft](#) for a case study-based analysis of the use of PKI for authentication within Microsoft):

- Certificate-based authentication can provide a high level of security for clients accessing the network by means of the following methods:
 - ▶ A VPN employing L2TP/IPSec with or without a RADIUS server such as Microsoft IAS
 - ▶ Wireless connections involving 802.1x authentication (see [Secure Wireless LAN Solution Architecture](#))
 - ▶ A secure Web-based connection using HTTPS with SSL or TLS
 - ▶ Smart card logon, both locally and across a VPN
- Authentication to Web sites using one of the methods supported by IIS 6.0 (anonymous, basic, digest, integrated Windows, or .NET Passport). For details on setting up client authentication to an intranet using IIS and Internet Explorer 6.0, refer to [HTTP-based Cross-Platform Authentication via the Negotiate Protocol](#).
- The need for authentication to or from networks operated by partner companies or trusted client companies. We discussed trust relationships for supporting this type of authentication earlier in this exam manual (see "Designing Forest and Domain Trust Models").
- Authentication of clients using various types of devices across a VPN. See [Enterprise Design for Remote Access](#) for further information.

- The use of advanced [two-factor authentication](#) technologies such as tokens or biometrics can provide a highly secure method of authenticating clients both locally and remotely.
- The need for authenticating Windows clients to non-Windows devices such as UNIX or main-frame computers. You may need additional services such as [Services for Unix, Host Integration Server](#), or other products.
- The need for authenticating non-Windows clients such as Macintosh or Linux computers, Palm Pilot PDAs, or various types of smart phones to your network.

Establishing Account and Password Security Requirements

We looked at the use of password policy and account lockout policy for establishing and enforcing account security earlier in this exam manual. When documenting and analyzing account and password security needs on your network, consider the following factors:

- *What account and password options are available for use with any type of client device connecting to your network from internal locations? Which options are still applicable of devices are connecting externally?*
- *How exactly do your account and password security policies match your company's written security policy? Third-party password filters are available that can enforce password requirements beyond those available in Windows password policy, such as avoiding the use of common English words or parts of the user name.*
- *What password capabilities are available with the various types of devices that might connect to your network? For example, a mobile phone may allow only the use of a 4-digit PIN. Try to avoid weakening your password policy to accommodate these devices. You may need some type of third-party solution to accommodate them.*

Designing a Security Strategy for Client Remote Access

We discussed designing remote access connectivity and authentication across a VPN earlier in this network, and showed you how to select authentication protocols for remote access including wireless access. In this section we look at designing remote access policies including access to internal resources and the use of IAS for authentication and accounting of remote access users.

Designing Remote Access Policies

[Remote access policies](#) provide a series of conditions that determine whether a user attempting remote access to your network is allowed to connect. Configured from the Routing and Remote Access snap-in or the Internet Authentication Service snap-in, remote access policies contain the following components (see *Figure 24*):

- *Conditions* - Attributes that a user's account must meet to satisfy the policy.
- *Permission* - If a user meets the conditions of the policy, you can specify that the user is either granted or denied remote access permission.
- *Profile* - Additional constraints that a user must meet for his connection to be accepted. Components included in the remote access profile include dial-in constraints, IP address and filters, use of multilink, authentication methods, encryption strengths, and advanced attributes.

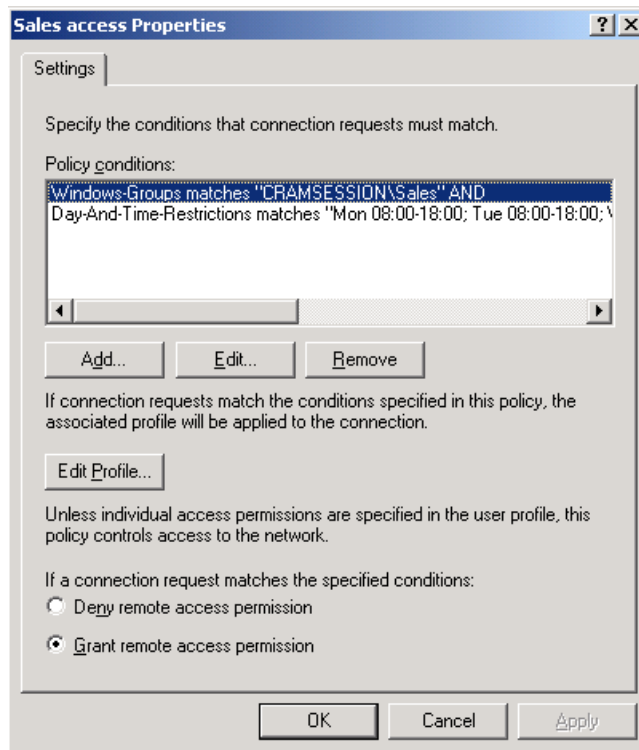


Figure 24 – You Can Configure the Conditions under Which a Remote Access Connection Is Accepted from a Remote Access Policy's Properties Dialog Box

The RRAS or IAS server evaluates each [connection attempt](#) according to the following multi-step procedure:

1. Each policy is evaluated in turn until the user meets the policy conditions. If she does not meet the conditions of the policy, the next policy in effect (as displayed in the RRAS snap-in) is evaluated. If the user does not meet the conditions of any available policy, she is denied access (implicit denial).
2. The Ignore-User-Dialin-Properties attribute, specified in the policy profile, is checked. If this attribute is set to **False**, the [dial-in setting](#) of the user's Properties dialog box in Active Directory Users and Computers is evaluated, as follows:
 - ▶ If it is set to **Deny**, the attempt is rejected.
 - ▶ If it is set to **Allow**, the attempt is evaluated against the policy profile (step 4) without regard to the permission setting on the policy properties.
 - ▶ If it is set to **Control Access Through Remote Access Policy**, the attempt is checked against the permission setting, and if this is set to Grant remote access permission, it is next checked against the policy profile.
3. If the Ignore-User-Dialin-Properties attribute is set to True, the remote access setting of the user's Properties dialog box is ignored, and the user is evaluated against the permission settings in the

- remote access policy, and if this is set to Grant remote access permission, it is next checked against the policy profile.
4. The conditions specified in the policy's profile are applied, and the user is granted access only if all conditions are met.

The following are several best practices you should employ when designing remote access policies (see [Best Remote Access practices](#) and [Dial-up remote access design considerations](#) for more information):

- Remember that remote access policies are evaluated in the order in which they appear in the details pane of the RRAS or IAS snap-in. When the conditions of a user attempting connection are met, policies lower in the sequence are ignored. Consequently, you should place the most specific policies first in this list.
- When users connect through multiple RRAS servers, use IAS for centralized authentication and authorization of remote access users. This enables you to configure the remote access policies once for all RRAS servers and reduces the chances of improper configuration when using more than one RRAS server. We discuss the use of IAS servers later in this section.
- For best client operating system security, require users that remotely access your network to use computers running Windows 2000, Windows XP, or Windows Server 2003. By default, computers running older operating systems do not support the latest security initiatives including authentication protocols and strong encryption levels.
- Use the [Connection Manager Administration Kit](#) (CMAK) to create customized connection profiles for your clients. We discuss CMAK in detail in the next section.
- Create one or more groups for users who need remote access to the network and configure remote access policies that place the appropriate conditions on these users in accordance with company security policies. You can then add users to or remove them from these groups as required to grant or revoke access.

Designing Access to Internal Resources

Once a remote access client has been authenticated to your network, the next question that comes to mind is "Is this client permitted to access all internal resources, or should it be restricted to only certain portions of the network?"

Besides the normal access controls such as shared folder and NTFS permissions, other access controls are available for controlling access by remote clients. Your design strategy should involve consideration of the following factors:

- Guidelines for security tools installed on clients requesting network access
- Use of managed VPN client connections
- Guidelines on the types of access that are allowed, as well as considerations about the level of access for different groups of remote access clients
- Use of tools such as Network Access Quarantine Control or Connection Manager

[Network Access Quarantine Control](#) is a new feature of Windows Server 2003 that limits the access available to an authenticated remote client until the configuration of the remote computer has been validated according to attributes included in a script written by the administrator. After a remote user is authenticated to the network, the client connection is quarantined while the script is run on the client computer.

Configuration items that can be checked on the client include the following:

- Operating system and service pack level in use
- Antivirus and antispyware software with up-to-date signature files
- Appropriate firewall software installed and activated on the Internet interface
- Routing to other networks disabled
- Password-protected screen saver operating

When the client passes these checks, it is fully authorized on the network and able to reach resources for which permissions are granted. If it fails these checks, it is directed to download sources for the required tools. Alternately, it can be given information for obtaining the proper tools and the connection then terminated.

Further information on Network Access Quarantine Control as well as a sample script can be obtained from Microsoft at [Network Access Quarantine Control in Windows Server 2003](#).

[Connection Manager](#) is a group of components that enables administrators to simplify the configuration of remote client computers including the creation of customized remote access profiles. It includes the Connection Manager Administration Kit (CMAK), which is a connection management utility and wizard that enables you to perform such tasks as distributing service profiles to remote clients, sending update information such as phone book lists for dialing to ISPs or branch offices to remote clients, and updating information to large numbers of remote clients. The following are some of the more important features that you can configure using CMAK:

- Phone book information, such as access numbers, used for VPN connection
- Dial-up networking entries including TCP/IP configuration information and basic and advanced security settings
- Routing table updates that define specific routes for network traffic accessing the VPN server
- Automatic configuration of Internet Explorer settings for a proxy server
- Custom actions such as programs that run automatically before, during, or after connection to the VPN
- Logon and phone book bitmaps and icons displayed to the user when connecting
- Notification area shortcut menu commands
- Help files and support information available to users
- License agreements and additional files used by the service profile

Using IAS to Provide Authentication and Accounting for Remote Network Access

[Internet Authentication Service](#) (IAS) is the Microsoft implementation of Remote Authentication Dial-In User Service ([RADIUS](#)), which is a centralized authentication, authorization, and accounting implementation standard described in Requests for Comments (RFCs) 2865 and 2866. It is a standard component of Windows Server 2003 that you can install from the Windows Components Wizard in Control Panel Add or Remove Programs.

The following are the major functions of IAS:

- IAS enables centralized authentication, authorization, and accounting for a series of RRAS servers, which operate as RADIUS clients. This provides the advantage of configuring remote access policies at a single location and consistent application of policies across all RRAS servers.
- IAS can also act as a [RADIUS proxy](#), which forwards authentication and accounting requests to other IAS servers. An ISP can operate a RADIUS proxy that forwards requests to IAS servers in multiple Active Directory forests.
- IAS can support 802.1x authentication for wireless clients.
- IAS can provide Network Access Quarantine Control, as discussed in the previous section.

[Designing the use of IAS](#) on your network can include one or more of the following activities:

- Inventorying the current network environment.
- Deciding on the number of IAS servers and their roles (RADIUS server, RADIUS proxy, as well as load balancing using multiple servers).
- Planning the RADIUS clients and the types of authentication to be used (We discussed available authentication protocols earlier in this exam manual.)
- [Optimizing](#) the IAS design including scaling the IAS servers for optimum performance, adding additional IAS servers as needed, implementing load balancing, optimizing RADIUS client performance, optimizing authentication and accounting, and optimizing IAS deployment in large organizations. This involves network design considerations such as the relative placement of IAS servers, IAS clients, domain controllers, and global catalog servers, as well as the possible deployment of IAS on a domain controller. The latter provides optimum authentication performance but may place an undesirable load on Active Directory service performance.
- Creation of a remote access policy strategy.
- Securing the remote access strategy, including the integration of IAS with certificate services and its *protection from server vulnerabilities*.
- *Implementing the IAS solution, including RADIUS servers, RADIUS proxies, and enabling compatibility with third-party remote access servers.*

Designing a Strategy for Securing Client Computers

Earlier in this exam manual, we looked at designing security for servers as a function of server role. Designing security for client computers involves similar strategies. Microsoft has published the [Windows XP Security Guide](#), which provides detailed information on securing computers running Windows XP Professional and operating in Active Directory domains requiring either normal security or high security, or those that function in a stand-alone environment (which includes those operating in a Windows NT 4.0 domain).

Designing a Strategy for Hardening Client Operating Systems

The first step in designing a strategy for hardening client computers is to design and implement an OU structure that takes into account the roles played by various client computers as well as the operating systems they run. In a similar manner to that previously shown for servers in *Figure 17*, you can create an OU named Clients with child OUs representing client computer roles such as the following:

- Desktop computers
- Laptop computers
- Kiosks
- Administrative workstations

Further, you can create OUs representing operating systems such as Windows 2000 Professional and Windows XP Professional. These OUs would be either parents or children of the role-based OUs, and would enable you to properly specify policies that apply to only a single operating system.

After you have created your OU design and placed the computer accounts in the appropriate OUs, you can design GPOs that apply to each situation, again in an analogous fashion to that previously described for servers.

As with servers, you can use [security templates](#) to apply security settings to all client computers subject to the GPO in which you apply them. You can also use Security Configuration and Analysis or the **secedit** command to apply security templates to client computers. This is especially useful in the case of stand-alone client computers.

Security templates enable you to specify security settings within the seven areas shown in *Figure 25*. The following are several recommendations you should follow when configuring security template-related settings for client computers:

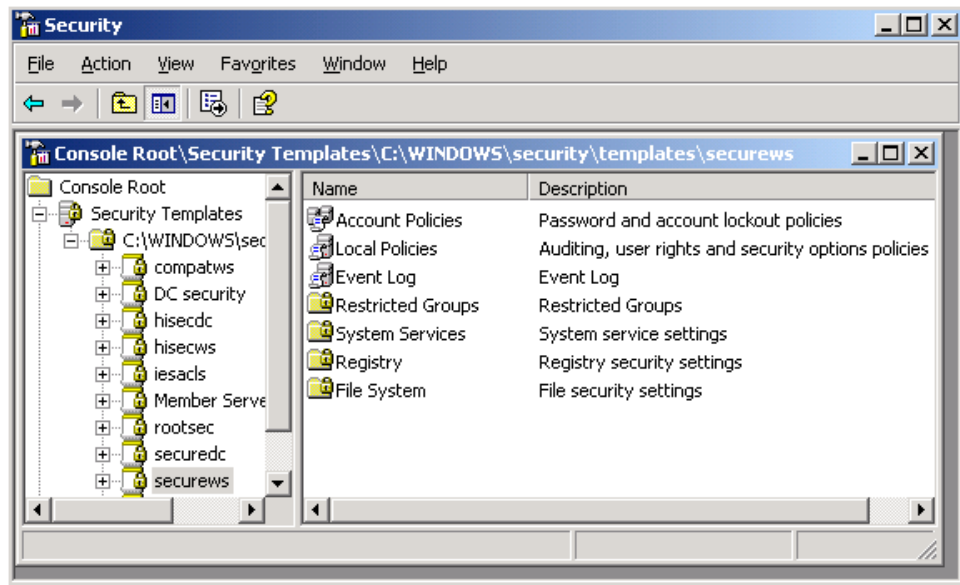


Figure 25 – You can Use Security Templates to Configure Security Settings for Client Computers

- *Account Policies* – Enables you to use an OU-based GPO to configure account policies that apply to users employing local user accounts for logging on to client computers. Remember that account policies that apply to domain user accounts must be configured in a domain-based GPO.
- *Local Policies* – Enables you to define auditing, user rights, and security options policies:
 - ▶ *Auditing* – While not as important as auditing on a server, client audit policies serve to track potential attempts at unauthorized access to client computers and determine actions that might have been attempted; for example, installation of Trojan horses or rootkits.
 - ▶ *User Rights* – Enables you to define the types of system activities users are permitted to perform; for example, accessing the computer locally or from the network.
 - ▶ *Security Options* – Enables you to define several logon-related options. Important ones you should consider include Do not display last user name, Message test for users attempting to log on, and Number of previous logons to cache (in case a domain controller is not available).
- *Event Log* – Enables you to define policies related to retention of the system, security, and application logs. This is crucial for ensuring that events logged during incursion attempts are maintained and available for viewing.
- *System Services* – Enables you to define the startup type for system services. You should disable services that are not required for normal operation such as Alerter, Computer Browser, Messenger, Net Meeting Remote Desktop Sharing, Universal Plug and Play Device Host, and Telnet, as well as IIS-related services.
- *Registry and File System* – Enable you to set permissions for folders, files, and registry keys.

See Chapter 3 of the Windows XP Security Guide (already referenced) and Chapters 2 to 7 of the [Threats and Countermeasures Guide](#) for detailed recommendations regarding each of these security settings. The Windows XP Security Guide also includes security templates that you can download and modify for each of the enterprise, high security, and stand-alone environments.

An important facet of client computer hardening not included in security templates is [software restriction policies](#), which enable you to specify which applications are permitted to execute on client computers. To define a software restriction policy, navigate to **Computer (or User) Configuration\Windows Settings\Security Settings\Software Restriction Policies**, right-click this node, and choose **New Software Restriction Policies**. The policy enables you to configure one of two available [security levels](#), and then define exceptions to the default rule:

- *Unrestricted* – Enables software to run according to the access rights of the logged on user. This is the default. If you specify this option, you next need to define exceptions that specify what types of software *are not* allowed to run.
- *Disallowed* – Prevents software from running, regardless of the access rights of the logged on user. If you specify this option, you next need to define exceptions that specify what types of software *are* allowed to run.

Having specified the desired security level, you can define exceptions according to the following available types of [policy rules](#):

- *Hash rule* – Defines allowable (or forbidden) applications according to the hash of the specified program, created by a standard hashing algorithm. When a user attempts to run a program, a hash is created and compared with the hash specified in the rule. If the hashes match, the program will be allowed or forbidden according to the defined rule.
- *Certificate rule* – Identifies allowable (or forbidden) applications according to the software's signing certificate. If the software is signed by a trusted publisher, it will be allowed or forbidden according to the defined rule.
- *Path rules* – Identify allowable (or forbidden) applications according to the path to the executable file or to [registry](#) keys used by the executable file. Four default path rules that allow system software to always be run are included (visible in the details pane of *Figure 26*). You should not modify these rules or the operating system will not work properly.
- *Internet zone rule* – Identifies allowable (or forbidden) applications according to the Internet zone from which Windows Installer packages have been downloaded. This rule applies only to Windows Installer packages.

To define software restriction rules, right-click **Additional Rules** and choose the appropriate rule type (*Figure 26*). Then provide the information requested by the dialog box that appears.

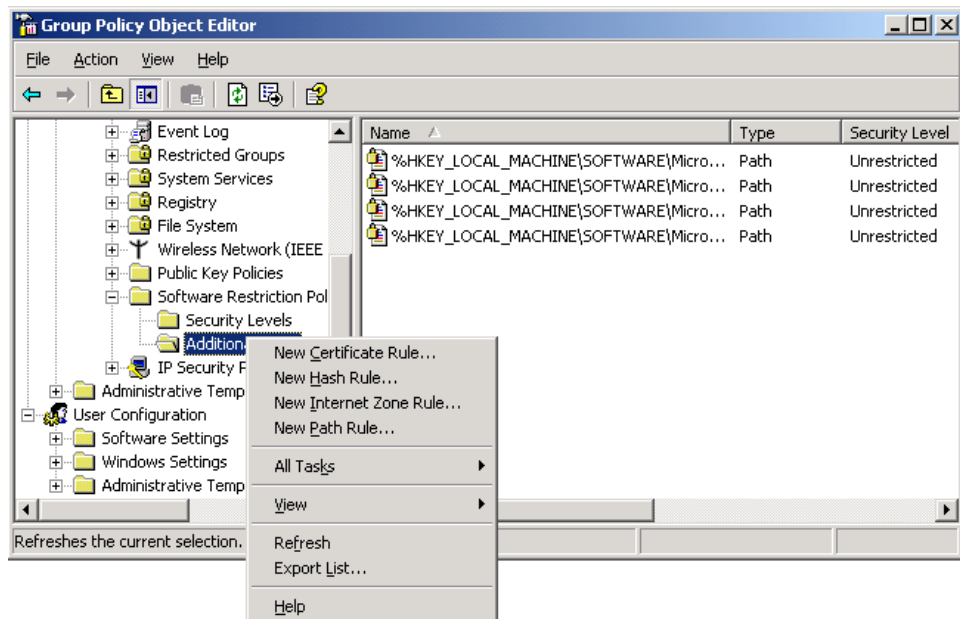


Figure 26 – You May Define Any of Four Types of Software Restriction Policy Rules

The following are several [best practices](#) you should follow when designing software restriction policies:

- *For high security environments, use a default rule of Disallowed* - This prevents all software from running, except where explicitly permitted. You should, however, use caution, and ensure that exceptions are specified for all applications that should be allowed to run on your network.
- *You can exempt administrators from software restriction policies if required* – Doing so would enable administrators to run network analysis or troubleshooting utilities in an environment where the default rule is Disallowed. To do so, right-click **Enforcement** and choose **Properties**. In the dialog box, select the **All users except local administrators** option.
- *Use a separate GPO to apply software restriction policies* – If you encounter problems, you can disable the GPO and then troubleshoot it. In particular, do not use the Default Domain Policy GPO for software restriction policies because you would need to disable all domain-based policies should problems occur with software restriction policies.
- *Use safe mode if you encounter problems* – Software restriction policies do not apply when you restart the computer in safe mode.
- *Use ACLs in conjunction with software restriction policies* – Users may attempt to move applications to bypass path rules. Use of ACLs can overcome this limitation of path rules.
- *Test your policies before implementing them* – This is particularly important if you have set a default policy of Disallowed, so that you can ensure that required applications will run.
- *Filter software restriction policies by security group membership* – This enables you to allow only members of specified groups to run certain applications.

Designing a Strategy for Restricting User Access to Operating

System Features All users are naturally curious about the tools and utilities provided with the default installation of operating systems such as Windows 2000 Professional and Windows XP Professional. They might poke around and change settings that prevent the computer from operating efficiently, or keep it from starting at all. You need to design policies that restrict users' ability to make configuration changes or access areas such as the Control Panel or Registry Editor.

Group Policy provides the [Administrative Templates](#) feature, available under both Computer Configuration and User Configuration, to restrict user access to portions of the operating system. Under Computer Configuration, it restricts all users of the computer to which the GPO is applied, and under User Configuration, it restricts access by users governed by the GPO to any computer that they might access. Administrative Templates are provided in the form of .adm files, which govern particular sets of Windows components. Additional .adm files are available for applications such as Office XP. You can add these files by right-clicking **Administrative Templates** and choosing **Add/Remove Templates**. Click **Add** and browse to the template file desired.

The Administrative Templates nodes of Group Policy provide hundreds of policies that you can use to limit user actions. See [Windows XP, Office XP, and Windows Server 2003 Administrative Templates](#) for a description of available policies including vulnerabilities addressed by each setting, recommendations for mitigating these vulnerabilities, and potential impacts of enabling the various settings. In addition, the Properties dialog box for each policy setting has an Explain tab (*Figure 27*) that assists you in determining which policy you should enable and the settings you should configure.

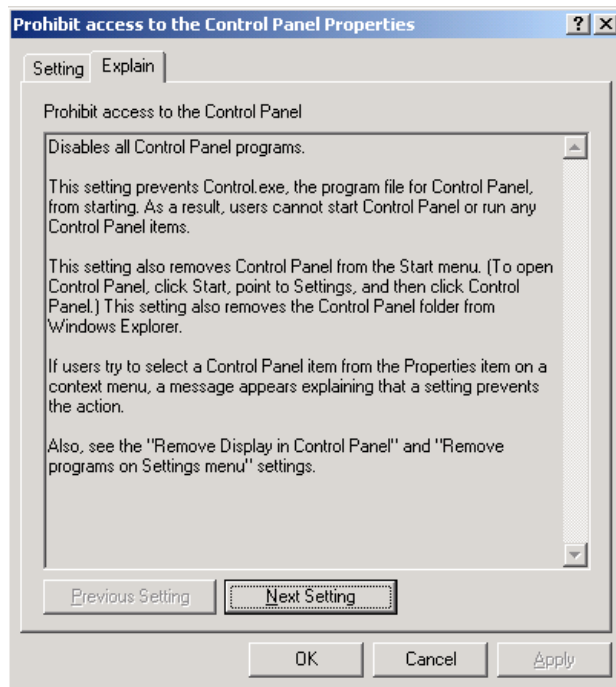


Figure 27 – The Explain Tab of an Administrative Templates Policy Provides Information on the Actions Taken If the Policy Is Enabled

Another means of controlling user access is by means of placing users in default groups provided by Windows that provide a limited set of privileges. The following are several [groups](#) you can use for this purpose:

- [Power Users](#) – Provides the user with a limited set of administrative capabilities such as the creation of local users and groups, running legacy applications, and sharing files.
- [Network Configuration Operators](#) – Users can manage several network configuration settings.
- [Print Operators](#) – Users can install, share, and manage printers attached to their own computers.
- [Remote Desktop Users](#) – Users can log on remotely to the desktop.

Keep in mind when designing policies for limiting user activity that other policies you can specify using features such as security templates and software restriction can also be used for this purpose. Refer to the links provided in this section for more information and guidance on selecting and enabling policies.

Practice Questions

Case Study 1

Background

You have been hired as a network consultant for SomeStore, Inc. SomeStore is a new and used book reseller with one main warehouse and multiple satellite storefronts with the same metropolitan area. All of the servers used by SomeStore have recently been upgraded to Windows Server 2003, with one loaded as an ISA server providing firewall protection to the corporate network. All of the remote sites are currently connected via dedicated dial-up lines (shown in Exhibit), using a single standalone DNS server in the main location for name resolution.

SomeStore has placed a wireless access point in each remote location so that clients of the in-store coffee shop can browse SomeStore's public online Web store using their own WiFi-compliant hardware. Transactional data from each location is updated to a central data repository in the main facility nightly.

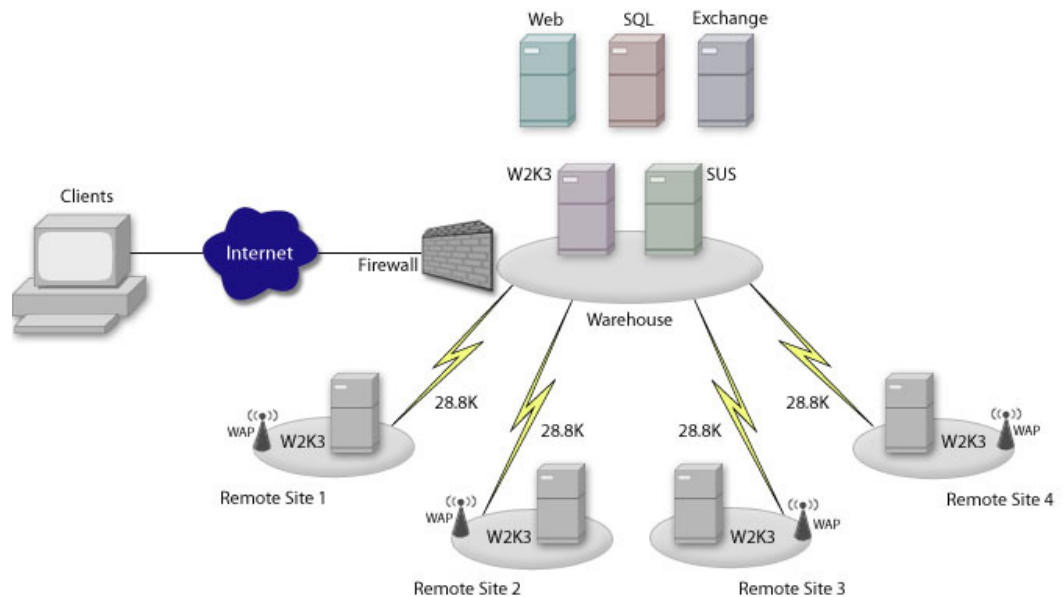
Interviews

Owner: I am concerned about storage of invoicing transactions at the remote sites. We need to store a week's work of billing; however, I'm told we have very little space available on those servers after we finish all of the upgrades our IT people have in mind.

Manager: I understand that our remote sites are occasionally experiencing problems with their name service due to traffic caused by transaction data uploads and some kind of information replication every 15 minutes. Remote users must be able to access our online storefront at all times, not so irregularly.

IT Manager: We will need to centralize auditing events in a central single location, as well as automate system updates across all locations. Constant replication is also causing some problems with DNS lookups and access to the online storefront. We are using separate subnets for each site to minimize cross talk, with a router/switch combination device providing isolation of each subnet accessed across the various modem links.

Exhibit(s)



1. You are directed to evaluate an upgrade to the remote sites' connectivity for high-speed connectivity to be used for open Internet access by coffee shop clients browsing wirelessly. What should you recommend?
Select the best answer.
 - A. Upgrade the dial-up links to dedicated leased line connectivity.
 - B. Replace the dedicated dial-up connections with on-demand dialing modems.
 - C. Replace all of the remote connections with public ISP-provided broadband connections.
 - D. Place only the wireless access point (WAP) connections on public ISP-provided broadband connections.

2. You need to minimize the impact of Active Directory replication between the remote locations and the main facility's server in order to facilitate the addition of more servers in each location in the future. How can you accomplish this task?
Select the best answer.
 - A. Place each location in its own domain.
 - B. Place each location in its own site.
 - C. Place each location in its own OU.
 - D. Manually modify each replication connection between servers.

3. You want to make sure that DNS access is available at all times from all locations, but do not want to purchase any additional hardware. You also must keep any update traffic to a minimum. How do you accomplish these tasks?

Select the best answer.

- A. Update the current name service zone to an Active Directory integrated zone and enable the DNS role on each of the remote servers.
- B. Update the current name service zone to an Active Directory integrated zone and enable the DNS role on each of the remote servers. You configure the replication cycle between zones to occur only on off-peak hours.
- C. Add the DNS role to each of the remote servers and configure replication to occur during off-peak hours.
- D. Configure the main location's DNS server with four stub zones, enabling the DNS service on each of the remote sites to support a Primary standalone zone for its own subnet alone.

4. You need to implement a patch management solution that allows updating of all domain members in all locations from a central approval site, but must limit bandwidth consumption to remote sites even if costs increase. How do you accomplish this?

Select the best answer.

- A. Test each hotfix and approve it from a central SUS server, utilizing SUS servers at each location to download updates from this server and automatically apply them to all client systems in their location.
- B. Test each hotfix and approve it on SUS servers located in each site, which will then apply them to each client system in their location.
- C. Test each hotfix and approve it on a SUS server in the main location. Configure all clients to automatically update using this server as their source.
- D. Test each hotfix and configure all clients to automatically update using the Windows Update service.

5. The IT Manager is concerned about potential threats posed to the remote site's servers by locally connected client systems using the WAP devices for network connectivity. How do you provide him with the ability to audit access attempts from the central office?

Select the best answer.

- A. Recommend placing each WAP device in a separate subnet from the servers in each location and configure the routers to monitor packet traffic between subnets.
- B. Recommend enabling MAC address control on each WAP device and configure the WAP equipment to log each failed access attempt to the main office's servers.
- C. Recommend enabling Wired Equivalent Privacy (WEP) on each WAP device and configure the WAP equipment to log each failed access attempt to the main office's servers.
- D. Recommend aggregating the audit logs from each of the remote servers using MOM.

6. You need to provide an automated mechanism for auditing the hardware and software inventory at all locations from the central site. What is the best method to accomplish this task? Select the best answer.

- A. Configure an instance of the SUS server at the main office.
- B. Configure an instance of the SUS server at the main office, as well as at the remote offices.
- C. Configure an instance of the SMS server at the main office.
- D. Configure an instance of the MOM server at the main office.
- E. Configure an instance of the MIIS server at the main office.

7. What backup strategy should you recommend for each remote site that will meet the requirements of the Owner? Each item represents a part of the complete solution. Select two.

- A. A full backup each week
- B. A full backup each day
- C. A daily backup each day
- D. An incremental backup each day
- E. A differential backup each day

Case Study 2

Background

PrepLogic is a publicly held corporation that operates a nationwide chain of department stores. Corporate headquarters is located in Raleigh, North Carolina. There is also a West Coast office located in San Francisco. The chain has 100 stores on the East Coast and 25 stores on the West Coast.

Network Infrastructure

PrepLogic's network consists of a single LAN at each store. Stores on the East Coast have a broadband connection to the Raleigh headquarters while West Coast stores are connected to the San Francisco office. There is a single broadband connection between the East Coast and West Coast, as illustrated in the Network Infrastructure exhibit.

All servers are running Windows 2000 Server. Client workstations run Windows 2000 Professional or Windows XP Professional. The Active Directory domain structure for PrepLogic is shown in the Active Directory Infrastructure exhibit.

Planned Changes

All Windows 2000 Professional systems will be upgraded to Windows XP Professional. All Windows 2000 Server systems will be upgraded to Windows Server 2003.

Interviews

Chief Executive Officer: One of the greatest challenges facing our business is the distance between our East Coast and West Coast operations. These difficulties hinder my ability to run the entire organization on a daily basis and impact shareholder value. Therefore, I plan to recommend to the Board of Directors that we divest ourselves of our West Coast operations and focus on building our core business on the East Coast.

Chief Financial Officer: I consider our relationship with Consulting Partners, Inc. a strategic asset to the company. We must be able to work closely with employees of Consulting Partners, and our IT infrastructure must be designed so that we can accept digital certificates issued by the Consulting Partners network.

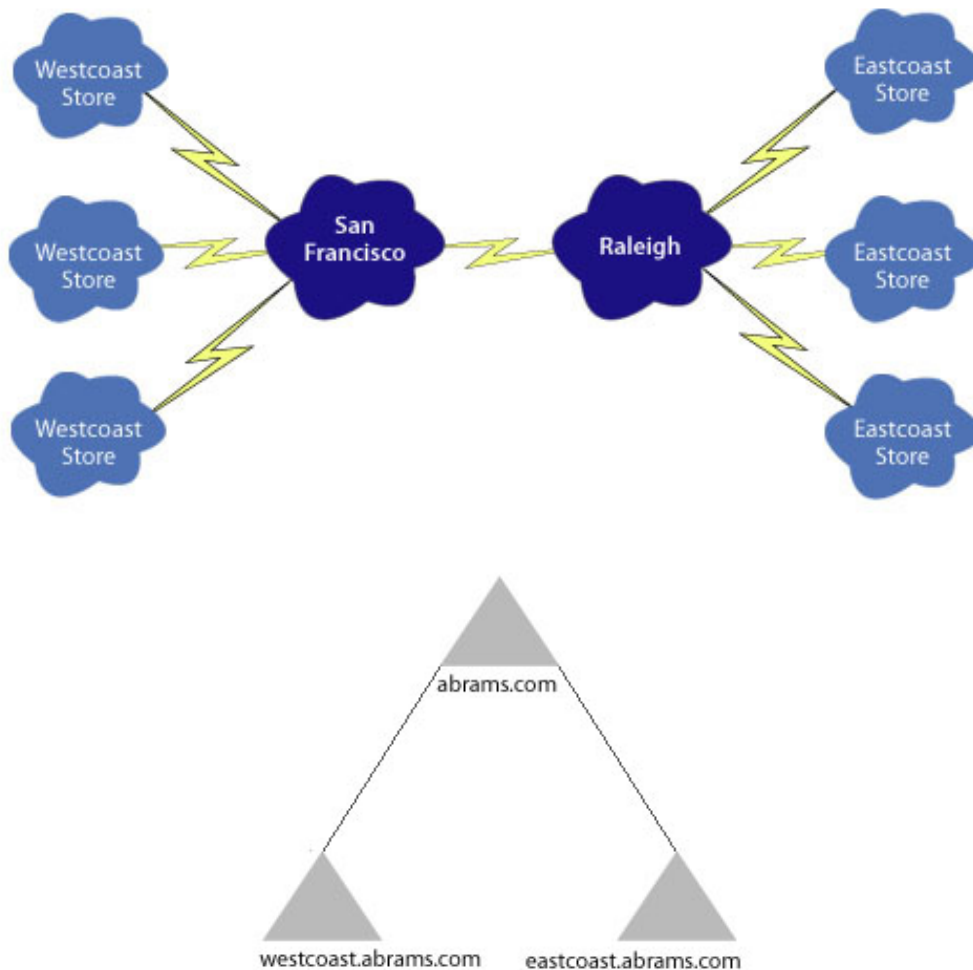
Chief Information Officer: We currently have a single WAN circuit connecting our West Coast and East Coast operations. This circuit is highly congested. We must work to minimize its utilization during business hours. I am also concerned about the patch level of systems on our network. I don't think that we have a handle on what systems need to be patched. I would like to see a report on the patches required by all PrepLogic systems.

West Coast IT Manager: I am constantly plagued by users who install software without permission. I need some way to control the applications that users run on their systems. However, I only have authority over the West Coast operations and can't impact operations on the East Coast.

Business Requirements

- The organization must be set up in such a manner that splitting the East Coast and West Coast assets is as simple as possible, in the event the business is sold.
- The certification authority (CA) hierarchy must be designed to ensure the security of the root CA's private key.
- The root CA must be located in the Raleigh office.
- Users and computers should be provided with public key infrastructure (PKI) certificates without manual intervention.
- The CA hierarchy should trust certificates issued by Consulting Partners but should ensure that this trust only applies to certificates from the Consulting Partners domain.

Exhibit(s)



1. You must create a PKI for PrepLogic. You need to choose a CA hierarchy that best meets the business requirements of the organization.
What should you do?
 - A. Implement an organizational CA hierarchy.
 - B. PKI should not be implemented at this time.
 - C. Implement a geographical CA hierarchy.
 - D. Implement a single CA on the East Coast.

2. You must design a security solution that meets the requirement expressed in the interviews for restricting software utilization on client workstations.
What should you do?
 - A. Create a software restriction policy and apply it to the PrepLogic.com domain.
 - B. Implement a Software Update Services (SUS) server in the PrepLogic.com domain and configure it to prohibit unauthorized software.
 - C. Implement a Software Update Services (SUS) server in the westcoast.PrepLogic.com domain and configure it to prohibit unauthorized software.
 - D. Create a software restriction policy and apply it to the westcoast.PrepLogic.com domain.

3. PrepLogic is in the midst of a transition from Windows 2000 systems to Windows 2003/XP systems. You must be aware of the security capabilities of each system on your network during the transition.
What should you do?
 - A. Know that Windows Management Instrumentation (WMI) filters will be ignored by Windows 2000 systems.
 - B. Know that Windows 2000 clients cannot connect to an SUS server.
 - C. Know that Windows 2000 clients will ignore any software restriction policies they encounter.
 - D. Know that Windows 2000 clients cannot use the Encapsulating Security Payload (ESP) functionality of Internet Protocol Security (IPSec) although they can use IPSec's Authentication Header (AH) technology.

4. You must design your CA hierarchy in such a manner as you provide the highest possible level of security to the root CA's private key.
What should you do?
 - A. Implement an enterprise root CA and take it offline.
 - B. Implement an enterprise intermediate CA and take it offline.
 - C. Implement a standalone root CA and take it offline.
 - D. Implement a standalone intermediate CA and take it offline.

5. You are designing a certificate distribution policy that must meet all business and security requirements.
What should you do?
- A. Create a custom certificate template for computers and configure it for automatic enrollment.
 - B. Use enrollment agents to facilitate the certificate distribution process.
 - C. Use Web-based enrollment through Internet Information Services (IIS) on a server other than the CA.
 - D. Use Web-based enrollment through Internet Information Services (IIS) on the CA.
6. You are designing the trust relationship between the PrepLogic CA and the Consulting Partners CA. You want to ensure that business requirements are met.
What should you do?
- A. Create a cross-certification with a basic constraint.
 - B. Create a cross-certification with a name constraint.
 - C. Create a cross-certification with an application constraint.
 - D. Create a cross-certification with a policy constraint.
7. You must prepare the report requested by the Chief Information Officer.
What should you do?
- A. Use the Resultant Set of Policy (RSOP) tool to create the report.
 - B. Use Group Policy Editor to create the report.
 - C. Use Software Update Services (SUS) to create the report.
 - D. Use Mbsacl.exe to create the report.

Case Study 3

Background

Goliath Construction is a nationwide builder of office buildings. Its corporate headquarters is located in Mineola, New York. Regional offices are located onsite for each new construction project undertaken around the country.

Network Infrastructure

Goliath's network infrastructure is quite simple. No public services are offered and a firewall prevents inbound connections. All server systems run Windows Server 2003. All desktop and laptop systems run Windows XP Professional.

Goliath's Active Directory infrastructure is shown in the Active Directory Infrastructure exhibit.

Interviews

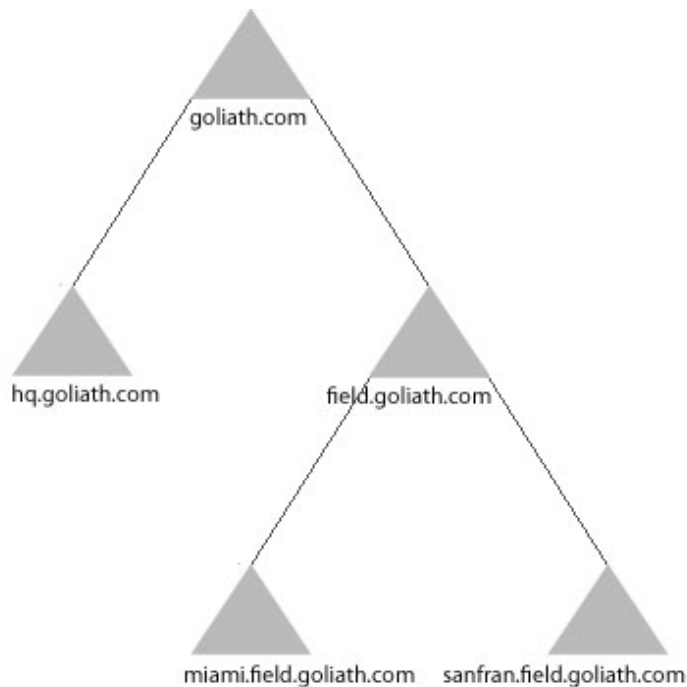
President and CEO: We're currently conducting two major building projects—one in Miami and one in San Francisco. We expect both of those projects to continue for five years. I hope to start three additional projects before the end of this year.

Chief Information Officer: We have a serious problem managing patches. I need a solution that enables me to check the patch status of every system on our network and applies critical updates as they're released by Microsoft.

Chief Security Officer: I'm concerned about the identity of users on our networks. We use wireless networks at our field locations, and I'm not sure the authentication is strong enough. I'd like to implement an organization-wide public key infrastructure (PKI) to provide digital certificate security throughout our enterprise.

Business Requirements

- Systems should be able to download critical updates from their local site.
- All critical updates must be approved by administrators in New York prior to use in the field.
- All certification authorities (CAs) should be standalone CAs.
- The Special Administration Console (SAC) must be available to provide administrators with a way to manage servers that fail to load the kernel properly.
- Users in the Miami domain should have more complex password security requirements than other users.

Exhibit(s)

1. You need to design a solution that meets the Chief Information Officer's requirements for checking system patch levels. You decide to use the Microsoft Baseline Security Assessment (MBSA) tool and must ensure that all systems are configured for scanning. What should you do?
 - A. Ensure that each system has the Remote Desktop Connection software installed.
 - B. Ensure that each system has Remote Registry enabled.
 - C. Ensure that either Remote Assistance or Terminal Services Client is installed on each client.
 - D. Ensure that the RestrictAnonymous registry key is set to 0 on each system.

2. You need to design the certificate distribution scheme for Goliath. Your solution must meet business requirements. What should you do?
 - A. Use manual enrollment.
 - B. Use automatic enrollment for user certificates.
 - C. Use automatic enrollment for computer certificates.
 - D. Use automatic enrollment for both user and computer certificates.

3. You need to design a solution that allows administrators to access servers when the kernel fails to load. Your solution must meet business requirements.

What should you do?

- A. Enable Remote Desktop for Administration (RDA).
 - B. Enable Terminal Services.
 - C. Enable Remote Assistance.
 - D. Enable Emergency Management Services.
4. You must design a strategy that allows for the more complex password security requirements specified by policy. Your solution must meet business requirements. Each answer presents part of the solution.

Choose two.

- A. Create a Group Policy object (GPO). Configure it with the password complexity requirements.
- B. Apply the GPO to the Default Domain Policy object for the Goliath.com domain.
- C. Apply the GPO to the miami.goliath.com domain.
- D. Enable complex password management on each system.
- E. Create a domain local group for all users that must have complex password requirements. Apply the GPO to that group.

5. You must design a Software Update Services (SUS) infrastructure to support the Chief Information Officer's requirements. How do you do so?

To answer, in the SUS Infrastructure exhibit, indicate where each SUS server should obtain its updates by dragging the appropriate choice from the accompanying list.

A. Update from Internet B. Update from HQ C. Update from Miami D. Update from San Francisco

Exhibit(s):



HQ SUS Server



Miami SUS Server



San Francisco SUS Server

Case Study 4

Background

The Mad Hatter is a chain of hat stores located in cities around the world. Corporate headquarters is located in London with national offices in New York and Tokyo.

Network Infrastructure

Mad Hatter's network infrastructure is shown in the Network Infrastructure exhibit. The Internet Information Services (IIS) server on the internal network provides services for internal clients only. All public services are outsourced to a third party. The Internet connection is used for outbound connections only.

Active Directory is currently in place. Domain Name System (DNS) services are integrated with Active Directory and Mad Hatter has an enterprise-wide public key infrastructure (PKI) that uses enterprise root and subordinate certification authorities (CAs) to provide certificates to all security principals on the network. All servers run Windows Server 2003. All client systems run Windows XP Professional.

Interviews


President: I am tired of synchronizing my laptop with the network. I want a wireless network installed in our corporate headquarters. This should not be very difficult -my mother installed one in her house and she doesn't know DOS from Windows!

Chief Information Officer: We can certainly install a wireless network to meet the president's requirements; however, we need to ensure that it has adequate security. Our rival firm, Bees in a Bonnet, is located in the same building, and I certainly don't want them snooping on our wireless network. We need to ensure that we have both strong authentication and encryption active.

IT Manager: My major concern is the patch level of systems under my control. I doubt that many of them are patched to the currently required level, and I'm loathe to invest the time necessary to bring them up to speed because they'll probably just fall out of date again anyway. I simply don't have the manpower to keep tabs on the critical updates that seem to come out almost weekly.

Business Requirements

- The wireless network must be encrypted and use strong authentication based upon our existing PKI.
- No funds are available for the purchase of additional hardware other than the wireless access point (WAP).
- Administrators must be able to resolve problems on client computers without physically traveling to remote sites.
- Only administrative computers should be able to connect to the database server via Telnet.
- All users should be able to ftp to the file server.
- No other Telnet connections are authorized.
- The network should use the minimum possible number of Internet Protocol Security (IPSec) filters.
- Nonrepudiation must be enforced for all e-mail messages

Exhibit(s)

1. You need to install a wireless access point on the network. You plan to make it a Remote Authentication Dial-In User Service (RADIUS) client and must select the appropriate RADIUS server. What should you do?
 - A. Use the IIS server as the RADIUS server.
 - B. Use the IAS server as the RADIUS server.
 - C. Use the file server as the RADIUS server.
 - D. Ensure that the database is ODBC compliant and use the database server as the RADIUS server.

2. You must select an authentication protocol for the organization's wireless network that meets all business requirements. What should you do?
 - A. Use Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication.
 - B. Use Protected Extensible Authentication Protocol-EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-EAP-MSCHAPv2) authentication.
 - C. Use smart card authentication.
 - D. Use Extensible Authentication Protocol-Message Digest 5 (EAP-MD5) authentication.

3. You are responsible for implementing security on the wireless network. You must select appropriate security measures that meet business requirements. Each answer presents a part of the solution. Choose three.
- A. Require validation of server certificates.
 - B. Implement Wired Equivalent Privacy (WEP) on the wireless network.
 - C. Ensure that certificate revocation checking is performed.
 - D. Implement Protected Extensible Authentication Protocol-Message Digest 5 (PEAP-MD5) authentication on as many systems as possible.
4. You need to design a solution that allows administrators to remotely administer client systems. Your solution must meet business requirements. What should you do?
- A. Enable Remote Desktop for Administration on client systems.
 - B. Enable Terminal Services on client systems.
 - C. Enable Remote Assistance on client systems.
 - D. Enable Remote Desktop Connection on the client systems.
5. You need to design a solution that ensures that the nonrepudiation business requirement is met. What should you do?
- A. Implement Encrypting File System (EFS).
 - B. Implement digital signatures.
 - C. Implement Secure Sockets Layer (SSL).
 - D. Implement access control lists (ACLs).
6. You need to develop a solution that meets the needs of the IT Manager. To answer, drag the appropriate steps to the right, in the proper order.
- | | |
|--|-------|
| A. Install SUS on the network. | _____ |
| B. Install the MSBA client on each system manually. | _____ |
| C. Use Software Update Services (SUS) to distribute the MBSA client to all systems | _____ |
| D. Configure a Group Policy object (GPO) that instructs all systems to obtain updates from the appropriate SUS server. | _____ |
| E. Use Microsoft Baseline Security Analyzer (MBSA) to generate a report detailing the current patch level. | _____ |

7. You need to design IPSec filters that meet the organization's business requirements. To answer, drag the appropriate policy setting to each empty block.

A. Allow B. Block C. No Policy

Exhibit(s):

| | Administrators Computers | All Systems |
|----------------------|-----------------------------|--------------------------|
| Telnet to Database | <input type="checkbox"/> | <input type="checkbox"/> |
| FTP to Database | <input type="checkbox"/> | <input type="checkbox"/> |
| Telnet to Fileserver | <input type="checkbox"/> | <input type="checkbox"/> |
| FTP to Fileserver | <input type="checkbox"/> | <input type="checkbox"/> |
| Telnet to any System | <input type="checkbox"/> | <input type="checkbox"/> |

Answers and Explanations

Case Study 1

1. Answer: D

Explanation A. Leased lines are more expensive than many other solutions. Simply upgrading the remote connectivity will still channel all data access through the main site's network, potentially swamping the net if a large number of clients are all browsing at the same time.

Explanation B. On-demand dialing may reduce connectivity time across the telephonic connections, but since because all lines are local, there is no per-minute charge for this service. On-demand dialing will require re-negotiation of the connection at least once every 15 minutes when an Active Directory replication cycle occurs. Constant renegotiation of the connection will impact connectivity negatively.

Explanation C. Without separating the WAP connections to the Internet or implementing an additional firewall at each location to protect the network's servers, this solution could allow many different types of attacks against the servers and other networked resources located in each location. It could while also preventing replication with the central site's server unless the main office firewall is modified to allow traffic from a public host.

Explanation D. This is the least costly and most secure solution. By removing the WAP connections to a separate public broadband connection, browsing clients can gain higher access to the Internet, while still being remaining able to access the online Web store without compromising the operational network supporting SomeStore's remote sites.

2. Answer: B

Explanation A. Although this solution might reduce some network traffic by reducing the number of changes that would have to be propagated throughout all domains in the forest, the cost of an additional server in each location to ensure adequate redundancy in each domain would be more than other solutions.

Explanation B. This solution will fulfill the requirement by allowing control over the replication cycle by allowing control of the replication between bridgehead servers in each site. This requires no additional hardware and minimal support effort so is the best solution.

Explanation C. Although this solution might prove useful in planning Group Policy application, it will do nothing to reduce replication traffic between sites.

Explanation D. This solution would provide the necessary control by modifying the default update cycle of each server's connection, but would require more effort than configuring bridgehead servers for each site in the event that more domain controllers were later added to any of the sites.

3. Answer: A

Explanation A. Active Directory integrated zones replicate with the normal Active Directory updates and so therefore require very little additional bandwidth.

Explanation B. Active Directory integrated zones replicate with the normal Active Directory updates, so they do not require the specification of a replication cycle.

Explanation C. Although this solution would ensure name service resolution at each site, it is not the least bandwidth-intensive solution. An Active Directory integrated zone would replicate with the normal Active Directory updates, providing a better solution.

Explanation D. The use of a stub zone allows delegation of name service lookup over a particular domain namespace, rather than a network subnet.

4. Answer: A

Explanation A. This is the best solution to the stated requirements.

Explanation B. This solution does not provide for a central approval and deployment solution, which would require the remote sites to be configured for automatic download and client update.

Explanation C. This solution will not minimize bandwidth consumption, as because each client must pull its own copy of the update from the main location's SUS server.

Explanation D. This solution will not minimize bandwidth consumption, as each client must pull its own copy of the update from the Windows Update site. This solution does not support control of patch approval.

5. Answer: D

Explanation A. This solution would require additional firewall solutions in each location, adding significantly to the cost. This is not the best solution to this question.

Explanation B. MAC address control restricts the devices that can connect to a WAP, but does not address the requirement of monitoring unauthorized access attempts from allowed client systems.

Explanation C. The WEP encryption scheme may serve to minimally improve security over data transmitted between clients and the network, but does nothing to provide monitoring of unauthorized access attempts from connected systems.

Explanation D. This is the best solution. The Microsoft Operations Manager server can aggregate event logs from all of the remote servers to a central repository and raise alerts if certain conditions are met.

6. Answer: C

Explanation A. This solution does not meet the requirements for auditing of software and hardware, only for automated update to the operating system.

Explanation B. This solution does not meet the requirements for auditing of software and hardware, only for automated update to the operating system using a decentralized distribution system.

Explanation C. This is the best solution. An SMS server can be used to poll all client systems and create an inventory of all hardware and software within the enterprise.

Explanation D. This solution does not meet the requirements for auditing client software and hardware, only for auditing events from the client systems.

Explanation E. The Microsoft Identity Integration Service does not support auditing of client software and hardware. It is used to allow authentication across multiple technologies and authentication domains.

7. Answers: A, D

Explanation A. This represents half of the proper solution. A full backup is required in order to facilitate the use of partial Incremental or Differential backups.

Explanation B. This solution would consume the largest amount of storage data for a week's backups, so it is not the correct solution.

Explanation C. A daily backup includes only those elements that have changed in the last 24 hours. In the event that backups are not made at the same time every day, this solution might not provide a complete backup.

Explanation D. This is half of the complete solution. In addition to a weekly full backup, daily Incremental backups will provide the least storage requirement for a full week's backup storage.

Explanation E. This is not the most efficient solution, as because each differential backup would include all changes since the last full backup. A daily Incremental backup would be more effective in terms of minimizing storage requirements.

Case Study 2

1. Answer: C

Explanation A. An organizational hierarchy is not well-suited to PrepLogic because it does not easily accommodate a split in the organization down the road.

Explanation B. Implementation of PKI is a business requirement. It is possible to choose a solution that balances the need for PKI with the requirement that the organization be prepared for divestiture of the West Coast assets.

Explanation C. A geographical CA hierarchy best facilitates the future sale of the West Coast assets while still allowing implementation of a PKI.

Explanation D. This solution does not facilitate the future sale of West Coast assets and fails to minimize WAN traffic.

2. Answer: D

Explanation A. This solution would apply to the entire domain and exceeds the authority of the West Coast IT Manager.

Explanation B. SUS is used to automate operating system updates and cannot control applications.

Explanation C. SUS is used to automate operating system updates and cannot control applications.

Explanation D. This solution addresses the business requirements expressed in the scenario.

3. Answers: A, C

Explanation A. WMI filters that appear in GPOs will be ignored by Windows 2000 clients.

Explanation B. Windows 2000 clients may connect to an SUS server.

Explanation C. Windows 2000 does not support software restriction policies and will ignore them when they appear in a GPO.

Explanation D. Windows 2000 clients can use both ESP and AH. More Information:

4. Answer: C

Explanation A. It is not possible to take an enterprise CA offline because it is integrated with Active Directory.

Explanation B. It is not beneficial to take an intermediate CA offline. Additionally, an enterprise CA may not be taken offline because it is integrated with Active Directory.

Explanation C. Taking the standalone root CA offline provides maximum security for the CA's private key.

Explanation D. It is not beneficial to take an intermediate CA offline.

5. Answer: A

Explanation A. Automatic enrollment meets the business requirement that certificates be issued without intervention.

Explanation B. Enrollment agents manually assist users with the creation of certificates. This does not meet the business requirement for certificate issuance without intervention.

Explanation C. The use of Web-based enrollment does not meet the business requirement for enrollment without manual intervention.

Explanation D. The use of Web-based enrollment does not meet the business requirement for enrollment without manual intervention.

6. Answer: B

Explanation A. Basic constraints are used to limit the path length for certificate chains.

Explanation B. Name constraints limit the namespace that may be trusted. In this case, a name constraint should limit the trust to certificates in the Consulting Partners namespace.

Explanation C. Application constraints limit the uses of trusted certificates.

Explanation D. Policy constraints limit the level of trust assigned to external certificates.

7. Answer: D

Explanation A. The RSoP tool analyzes the impact of applying multiple Group Policy objects (GPOs) to the same entity.

Explanation B. The Group Policy Editor is used to manage Group Policy objects (GPOs).

Explanation C. SUS distributes critical updates to systems on the network but does not have a scanning or reporting function.

Explanation D. Microsoft Baseline Security Analyzer determines the patches that need to be installed on a system.

Case Study 3**1. Answer: B**

Explanation A. Remote Desktop Connection is not necessary for MBSA scanning.

Explanation B. MBSA uses the Remote Registry service to assess the system's patch level. If this service is not available, MBSA will not be able to scan the system.

Explanation C. Remote Assistance and Terminal Services Client are remote access solutions and are not necessary for MBSA to function.

Explanation D. The RestrictAnonymous registry key is used to control anonymous access to directory services but is not used by MBSA.

2. Answer: A

Explanation A. You must use manual enrollment, as it is the only certificate distribution method supported by standalone CAs.

Explanation B. Automatic enrollment is not available for standalone CAs.

Explanation C. Automatic enrollment is not available for standalone CAs.

Explanation D. Automatic enrollment is not available for standalone CAs.

3. Answer: D

Explanation A. RDA will only function if the kernel loads.

Explanation B. Terminal Services will only function if the kernel loads.

Explanation C. Remote Assistance will not function if the kernel fails to load. Furthermore, it is not available on server operating systems.

Explanation D. Emergency Management Services allows use of SAC when the kernel fails to load.

4. Answers: A, C


Explanation A. You should create a GPO with the complexity requirements.

Explanation B. These requirements should apply only to the Miami domain. Therefore, the GPO should not be applied to the root domain.

Explanation C. The GPO should be applied only to the Miami domain. **Explanation D.** It is not necessary to enable complex password management on each system.

Explanation E. The password complexity requirements are not designed for a specific group of users, rather they are for an entire domain.

5. Answer:



Explanation: The correct answer is to have the HQ SUS server update from the Internet and all other servers update from the HQ site.

Case Study 4

1. Answer: B

Explanation A. RADIUS services are provided by Microsoft Internet Authentication Services (IAS).

Explanation B. RADIUS services are provided by Microsoft Internet Authentication Services (IAS).

Explanation C. RADIUS services are provided by Microsoft Internet Authentication Services (IAS).

Explanation D. RADIUS services are provided by Microsoft Internet Authentication Services (IAS).

2. Answer: A

Explanation A. EAP-TLS authentication requires an existing PKI and provides the most secure level of authentication.

Explanation B. PEAP-EAP-MSCHAPv2 does not take advantage of the organization's existing PKI.

Explanation C. Smart card authentication would require the acquisition of additional hardware resources and does not meet business requirements.

Explanation D. EAP-MD5 authentication is inherently insecure and does not provide mutual authentication.

3. Answers: A, B, C

Explanation A. Mutual validation of certificates protects wireless network participants from man-in-the-middle attacks.

Explanation B. WEP provides confidentiality for data while in transit.

Explanation C. Certificate revocation checking ensures that stolen systems may not be used on the network between the time of theft and certificate expiration.

Explanation D. PEAP-MD5 is inherently insecure and would compromise the security of the wireless network.

4. Answer: C

Explanation A. Remote Desktop for Administration is a tool used to administer Windows Server 2003 systems and is not available on Windows XP Professional systems.

Explanation B. Terminal Services is not available on Windows XP Professional systems.

Explanation C. Remote Assistance allows administrators to provide support to client systems without physically traveling to the system.

Explanation D. Remote Desktop Connection is used by the administrator to make the connection and should not be installed on the client systems.

5. Answer: B

Explanation A. EFS provides file system confidentiality but does not enforce nonrepudiation.

Explanation B. Digital signatures may be used to enforce nonrepudiation.

Explanation C. SSL may be used for authentication and confidentiality but does not enforce nonrepudiation.

Explanation D. ACLs may be used to enforce file permission settings but not nonrepudiation.

6. Answer:

- A. Install SUS on the network.
- B. Install the MSBA client on each system manually.
- C. Use Software Update Services (SUS) to distribute the MBSA client to all systems
- D. Configure a Group Policy object (GPO) that instructs all systems to obtain updates from the appropriate SUS server.
- E. Use Microsoft Baseline Security Analyzer (MBSA) to generate a report detailing the current patch level.

E.
A.
D.

Explanation: The MBSA report provides the initial analysis requested by the IT Manager. SUS is necessary to meet the requirement for automatic distribution of updates. A GPO should be used to ensure that all systems obtain updates from the appropriate SUS server.

No client is necessary for MBSA to scan systems. Target systems must have the Remote Registry enabled, however.

7. Answer:

| | Administrators Computers | All Systems |
|----------------------|-----------------------------|-----------------------------|
| Telnet to Database | <input type="checkbox"/> A. | <input type="checkbox"/> B. |
| FTP to Database | <input type="checkbox"/> C. | <input type="checkbox"/> C. |
| Telnet to Fileserver | <input type="checkbox"/> C. | <input type="checkbox"/> B. |
| FTP to Fileserver | <input type="checkbox"/> C. | <input type="checkbox"/> B. |
| Telnet to any System | <input type="checkbox"/> C. | <input type="checkbox"/> C. |