

Microsoft Server 2003

Active Directory and Network Infrastructure

(70-297)

Microsoft Certified
Systems Engineer (MCSE)



**Smarter
Training**

This LearnSmart exam manual covers the most important topics with which you must be familiar in order to successfully complete the Server 2003 Active Directory and Network Infrastructure exam (70-297). By studying this exam manual, you will master an array of exam-related material, including:

- Creating the Conceptual Design by Gathering and Analyzing Business and Technical Requirements
- Creating the Logical Design for an Active Directory Infrastructure
- Creating the Physical Design for an Active Directory and Network Infrastructure
- 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 Infrastructure (70-297) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC
Product ID: 2182
Production Date: July 18, 2011
Total Questions: 30

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	7
What to Know	7
Tips	7
Analyzing the Existing DNS Implementation	8
Naming Conventions	9
<i>Distinguished Names</i>	9
<i>Relative Distinguished Names</i>	10
<i>User Principal Names</i>	10
<i>Globally Unique Identifiers</i>	10
Analyze DNS for Active Directory Directory Service Implementation.....	10
Creating the Logical DNS Design	11
Design a DNS Name Resolution Strategy	11
<i>Namespace Design</i>	11
<i>DNS Interoperability</i>	11
<i>Zone Requirements</i>	13
<i>DNS Security</i>	15
Creating the Physical DNS Design.....	17
Design DNS Server Placement	17
Domain Controllers	18
Flexible Single Master Operation (FSMO) Roles	18
<i>Forest Wide Operation Master</i>	18
<i>Domainwide Operation Masters</i>	19
Placement of FSMO Roles.....	20
Design the Active Directory Infrastructure	
to Meet Business and Technical Requirements	21
Design the Envisioned Administrative Model	21
Create the Conceptual Design of the Active Directory Forest Structure	22
Create the Conceptual Design of the Active Directory Domain Structure	23
<i>Trust Relationships</i>	24
Create the Conceptual Design of the Organization Unit (OU) Structure	26
Design the Network Services Infrastructure to Meet Business and	
Technical Requirements.....	27
Create the Conceptual Design of the DNS Infrastructure	27

Create the Conceptual Design of the WINS Infrastructure	27
Create the Conceptual Design of the DHCP Infrastructure	27
Create the Conceptual Design of the Remote Access Infrastructure	28
Design an OU Structure	29
Identify the Group Policy Requirements for the OU Structure	29
Design an OU Structure for the Purpose of Delegation	29
Design a Security Group Strategy	29
Define the Scope of a Security Group to Meet Requirements	30
<i>Changing Group Scopes</i>	30
<i>Nesting</i>	30
Define Resources Access Requirements	31
Define Administrative Access Requirements	31
Define User Roles	31
Design an Active Directory Site Topology	32
Design Sites	32
<i>Intra-Site Replication</i>	33
<i>Intersite Replication</i>	34
Identify Site Links	34
<i>Replication Interval</i>	34
<i>Replication Schedule</i>	35
<i>Replication Transport</i>	35
<i>Link Cost</i>	35
New Features in Windows Server 2003	36
Active Directory	36
Security	36
Analyze the Impact of Active Directory on the Existing Technical Environment	37
Analyze Hardware and Software Requirements	37
Analyze Interoperability Requirements	37
Analyze Current Level of Service within an Existing Technical Environment	38
Analyze Existing Network Operating System Implementation	38
Identify the Existing Domain Model	38
Identify the Number and Location of Domain Controllers on the Network	40
Identify the Configuration Details of all Servers on the Network	40

Analyze Security Requirements for the Active Directory Service	40
Analyze Current Security Policies, Standards, and Procedures	40
Identify the Impact of Active Directory on the Current Security Infrastructure	41
Identify the Existing Trust Relationships	41
Identify Network Topology and Performance Levels	43
Identify Constraints in the Current Network Infrastructure	43
Interpret Current Baseline Performance Requirements for Each Major Subsystem	44
Analyze the Impact of the Infrastructure	
Design on the Existing Technical Environment	45
Analyze Hardware and Software Requirements	45
Analyze Interoperability Requirements	46
Analyze Current Level of Service within the Existing Technical Environment	46
Analyze Network Requirements	46
Design a Computer and User Authentication Strategy	47
Identify Common Authentication Requirements	47
Select Authentication Mechanisms	47
<i>Kerberos</i>	48
<i>PKI</i>	49
<i>Smart Cards</i>	49
Design a User and Computer Account Strategy	50
Specify Account Policy Requirements	50
Design Migration Paths to Active Directory	51
Define Whether Migrations will Include an In-place Upgrade, Domain Restructuring, or Migration to a New Active Directory Environment	52
Design a Strategy for Group Policy Implementation	53
Design the Administration of Group Policy Objects	53
Design the Deployment Strategy of GPOs	53
Design a NetBIOS Name Resolution Strategy	54
Design a WINS Replication Strategy	54
Design a Remote Access Strategy	56
Specify the Remote Access Method	56
Specify the Authentication Method for Remote Access	56
Design the Remote Access Infrastructure	57
Plan Capacity	57

Design Security for Remote Access Users	58
Identify Security Host Requirements	58
Identify the Authentication and Accounting Provider	58
Design Remote Access Policies	58
<i>Conditions</i>	58
<i>Permissions</i>	58
<i>Profile</i>	59
<i>Remote Access Policy Evaluation</i>	59
Specify Logging and Auditing Settings	59
Design a DNS Service Implementation	60
Design a Strategy for DNS Zone Storage	60
Specify the Use of DNS Server Options	60
Identify the Registration Requirements of Specific DNS Records	61
Specify the Server Specifications to Meet System Requirements	62
<i>Web Server</i>	62
<i>Standard Server</i>	62
<i>Enterprise Server</i>	63
<i>Datacenter Server</i>	63
<i>Hardware Requirements</i>	64
Design Internet Connectivity	64
Design a Network and Routing Topology for a Company	65
Design a TCP/IP Addressing Scheme Through the Use of IP Subnets	65
Specify the Placement of Routers	66
Design IP Address Assignment by Using DHCP	66
Design a Perimeter Network	67
Practice Questions	68
Answers and Explanations	78

Abstract

This Exam Manual will help you prepare for Microsoft exam 70-297, Designing a Microsoft Windows Server 2003 Active Directory and Network Infrastructure. Exam topics cover the official objective domains: Creating the Conceptual Design by Gathering and Analyzing Business and Technical Requirements, Creating the Logical Design for an Active Directory Infrastructure, Creating the Logical Design for a Network Services Infrastructure, and Creating the Physical Design for an Active Directory and Network Infrastructure.

What to Know

The official domains of the exam are:

- Analyzing for Active Directory
- Identifying business requirements
- Designing an Active Directory Infrastructure
- Designing a network services infrastructure
- Designing a naming strategy
- Designing a name resolution strategy
- Migration strategies
- Active Directory site topologies
- Designing a remote access strategy
- Server placement
- Internet connectivity
- Security

Tips

Exam 70-297 consists of 112 questions that cover a variety of topics related to Active Directory and network infrastructure design. You can expect to see questions on Domain Name System (DNS) namespace, Active Directory design considerations, server placement, Group Policy, organizational units (OUs), and network services, such as Remote Access Service (RAS) and Dynamic Host Configuration Protocol (DHCP).

The exam is made up of 11 different case studies, each followed by a series of questions that require knowledge of Active Directory and network infrastructure design. Although Microsoft covers several domains on the exam, some domains receive much more coverage than others.

One good way to prepare for the exam is to spend some time running Windows Server 2003 on your own test network. One of the most important keys to success on any certification exam—especially a Microsoft exam—is hands-on experience. Install, configure, and practice using the software on which you will be tested—Windows Server 2003 in this case—and the various network services included with it.

You'll have plenty of time to complete the test. Don't rush it. The exam allows you to mark questions you are unsure of and return to them later. Use this feature to your advantage. On the first pass, work through the exam, answering all the questions that you are absolutely sure of. On the second pass, spend time on the questions that present 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.

Analyzing the Existing DNS Implementation

When Active Directory was introduced with Windows 2000, it incorporated the DNS naming convention. Windows Server 2003 also uses the DNS naming convention as its primary naming convention.

A **namespace** can be defined as any bounded area in which a name can be resolved. Windows Server 2003 has adopted the DNS namespace to remain interoperable with Internet technologies and standards. This means that DNS is used to resolve Internet names and used on internal networks to resolve and locate computers.

DNS names are much simpler to remember than Internet Protocol (IP) addresses. It allows users to connect to servers using the same naming convention as they would if connecting to servers on the Internet. Additionally, DNS names tend to be more static than IP addresses, especially if DHCP is implemented.

The DNS namespace is hierarchical in structure, unlike the Network Basic Input/Output System (NetBIOS) naming convention which is a flat namespace. Each domain with Active Directory must have a DNS name. The DNS name of a domain represents its position within the hierarchy. If a child domain is added to the hierarchy, it inherits a portion of its namespace from its parent domain.

The hierarchical naming structure in Active Directory begins with a root domain and can potentially have second-level domains underneath it that will share part of the root domain's namespace. When looking at the DNS structure of Active Directory, there are two types of namespaces: **contiguous** and **disjoint**.

With a contiguous namespace, the child object in a hierarchy inherits a portion of its namespace from its parent domain. With a disjoint namespace, a child object's name is independent of its parent domain. A single forest can have both contiguous and disjoint namespaces.

Let's look at an example to illustrate the difference between the two namespaces. Suppose you create a new Active Directory domain named `cramsession.com` as shown in *Figure 1*. This domain becomes the forest root from which any other domains will inherit a portion of their namespace if they are added to the existing tree.

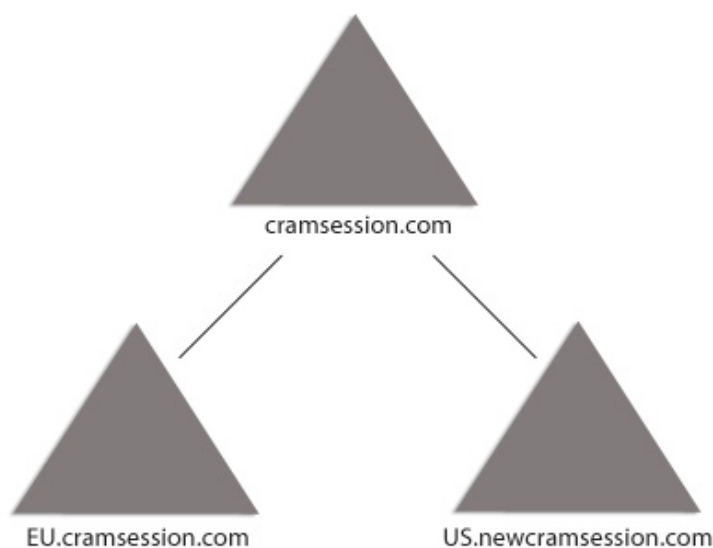


Figure 1 – Contiguous Namespace

On the other hand, as you will see when you get into the design aspects of Active Directory, there may be instances when you do not want a new domain to share the existing namespace but still want it within the existing forest. In this situation, you can create a new tree within the existing forest, such as newcram-session.com, thereby creating a disjointed namespace, as shown in *Figure 2*.

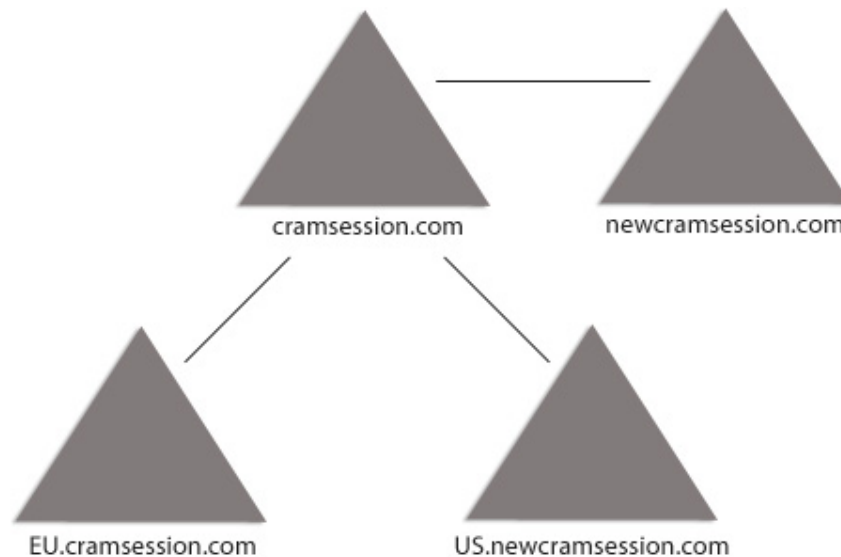


Figure 2 – Disjointed Namespace

See [How Active Directory Searches Work](#) for more information.

Naming Conventions

Active Directory supports a variety of different naming conventions, such as distinguished names, relative distinguished names, user principal names, and globally unique identifiers. Each object you create within Active Directory can be identified by one of these naming conventions.

Distinguished Names

Each object that is created within Active Directory is assigned a distinguished name (DN). It uniquely identifies the object as well as its location within an Active Directory forest. For example, let's say you create a user account within Active Directory for a user named Sean Massa. The DN for the user account may look similar to the following:

CN=Sean Massa,OU=SALES,DC=BAYSIDE,DC=COM

- The DC attribute identifies the domain which in this case is BAYSIDE.COM.
- The OU attribute identifies the complete path to the location of the object within the Active Directory forest.

- The object is located in the SALES OU, which stores user accounts for members of the sales department.
- The CN attribute represents the user's first and last names.

All objects within an Active Directory forest must have a unique DN.

See [Object names](#) for more information.

Relative Distinguished Names

Searching for an object by its DN can be difficult, especially if you don't know its exact location. The relative distinguished name (RDN) is part of the name that is an actual attribute of the object and it allows you to search for an object based on one of its attributes.

The RDN is part of the DN that uniquely identifies the object within its container. In the example in the preceding section, the RDN for the user account would be Sean Massa. No two objects within the same container can have the same RDN; therefore, you cannot have two Sean Massa accounts within the Sales OU. However, objects within the hierarchy can have the same RDN if they are in different containers because they would have unique DNs.

User Principal Names

Because the syntax of DNs and the actual names themselves can be difficult to remember, you can also use user principal names (UPNs), which are shorter than DNs and easier to remember. The UPN is usually made up of a shorthand name for the user followed by the domain name in which the account resides. For example, the user Sean Massa has a user account in the BAYSIDE domain. The UPN for the account may look something similar to SMassa@bayside.com.

Globally Unique Identifiers

Every object created within Active Directory is assigned a globally unique identifier (GUID) that becomes the object's permanent identifier. A GUID is a 128-bit number that is guaranteed to be unique within an Active Directory forest. The GUID assigned to an object never changes, even if the object is renamed or moved. This is unlike DNs and RDNs that change when an object is moved or renamed.

See [How Active Directory Searches Work](#) for more information.

Analyze DNS for Active Directory Directory Service Implementation

The existing DNS namespace impacts the Active Directory design. Some organizations implement the same namespaces both internally and externally. Conversely, others choose to implement DNS namespaces that are completely different—the internal and external namespaces are unique. Finally, the internal namespace may be a subdomain of the external namespace. Document the namespaces for later reference.

Creating the Logical DNS Design

Design a DNS Name Resolution Strategy

Designing a DNS name resolution strategy includes several steps. You must decide on the namespace design that best meets the business requirements. Identify any interoperability requirements with Active Directory, Windows Internet Naming Service (WINS), DHCP, and UNIX. The DNS name resolution strategy must also meet security requirements.

When selecting a DNS namespace to use, register a unique DNS name.

Namespace Design

- External and internal DNS namespaces can be the same.
 - Ensure that resource records from the internal DNS servers are not replicated to the external DNS servers.
 - This design option is the least secure.
- External and internal DNS namespaces can be different.
 - The combination provides the highest level of security.
 - The internal namespace is not exposed to the Internet.
- Delegate a subdomain for the internal DNS namespace.
 - This design option is more secure than using identical external and internal namespaces.

See [Namespace planning for DNS](#) for more information.

DNS Interoperability

- DNS must interoperate with many other network services, such as DHCP, WINS, and Active Directory.
- Windows Server 2003 DNS can interoperate with UNIX Berkeley Internet Name Domain (BIND).
- DNS servers have to support SRV records to interoperate with Active Directory.
 - SRV records identify servers on the network providing specific services such as Active Directory.
- UNIX BIND versions earlier than 4.9.7 do not support SRV records.
- UNIX BIND version 8.2.1 supports SRV records as well as secure updates and incremental zone transfers.
- WINS is needed only if there are clients or applications on the network that use NetBIOS.
 - If a name cannot be resolved using the DNS namespace, a DNS server can query a WINS server.

DNS can interoperate with a DHCP server. Dynamic DNS was introduced in Windows 2000 and is included in Windows Server 2003. With dynamic updates, the DNS server allows DHCP servers and DHCP clients to dynamically register A records and pointer (PTR) records. Windows 2000, Windows XP, and Windows Server 2003 clients are capable of updating their own records. For clients that cannot perform this function, a DHCP server can perform the updates on their behalf. For example, when a Windows 95 client leases an IP address, the DHCP server can update the client's records with the DNS server on its behalf.

By default, Windows 2000 clients and later are configured to update their own A records with the DNS server and the DHCP server updates the PTR records. Alternatively, you can configure the DHCP server to also update A records on their behalf. For pre-Windows 2000 clients, the DHCP server updates both A and PTR records. This is enabled on the DHCP server, as shown in *Figure 3*.

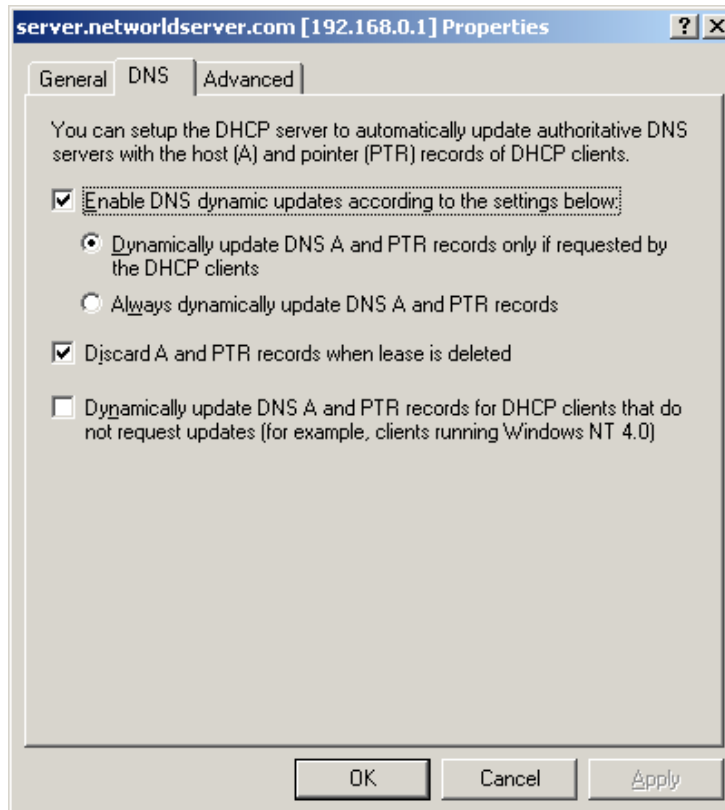


Figure 3 – DHCP/DNS Integration in the Server's Properties Dialog Box

See [Deploying Domain Name System \(DNS\)](#) for more information.

Zone Requirements

A zone file stores resource records for a domain. *Table 1* lists some of the common types of resource records.

Resource Record	Description
Start of Authority (SOA)	Located at the beginning of each zone file.
SRV record	Identifies servers on the network that are running specific services.
Host Address (A) Record	Maps a DNS name to an IP address.
Mail Exchanger (MX) Record	Routes messages to a specified mail exchanger for a specified DNS domain name.
Pointer (PTR) Records	Points to a location in the DNS namespace. PTR records are normally used for reverse lookups.
Alias (CNAME) Record	Specifies another DNS domain name for a name that is already referenced in another resource record.

Table 1 – Common Types of Resource Records

Windows Server 2003 DNS supports different [types of zone files](#), as shown in *Figure 4*. The type of zone file you choose depends on the business requirements.

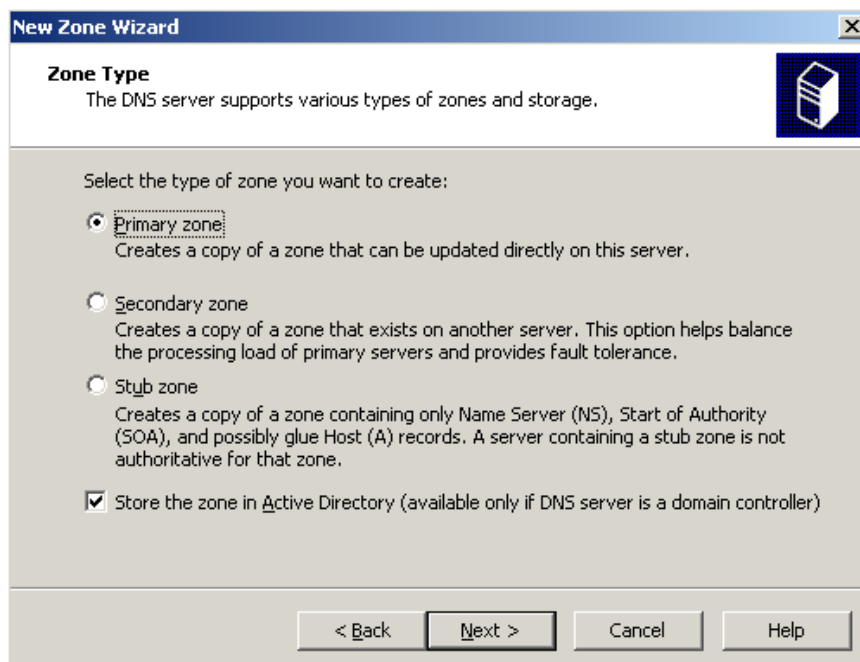


Figure 4 – DNS Zone Types

- *Primary zone:*
 - This is the working copy of the zone file.
 - Changes and updates must be made to the primary zone file.
 - A primary zone is normally used when there are older DNS systems on the network or UNIX DNS.
- *Secondary zone:*
 - This is a read-only copy of a DNS zone.
 - Updates are replicated from a primary zone or another secondary zone.
 - A secondary zone provides fault tolerance and load balancing.
- *Stub zone:*
 - This type of zone is used to identify authoritative DNS servers.
 - It only has the resource records needed to identify which DNS servers are authoritative for a zone.
 - Facilitates name resolution between separate namespaces.
- *Active Directory-integrated zone:*
 - Zone information is stored within Active Directory.
 - This zone must be stored on a domain controller.
 - An Active Directory-integrated zone increases fault tolerance because zone data is stored on all domain controllers.
 - This is the recommended zone type.

When zone information is stored within Active Directory, zone updates are replicated differently than in a standard primary/secondary scenario. DNS notification is no longer needed and configuring a notify list is unnecessary. Instead, DNS information replicates every time Active Directory replicates.

Multimaster replication is used when zones are stored in Active Directory. Normally, DNS updates are performed using a single master update model. This means that updates occur on a single DNS server that is authoritative for the domain and are then replicated to secondary DNS servers. When zone information is stored within Active Directory, a multimaster update model is used and any DNS server acting as a domain controller becomes authoritative for the zone. Updates can then be performed on any domain controller running the DNS service. This also eliminates the primary DNS server as being a single point of failure, which is the case in the single master update model.

Zone transfers are used to replicate zone updates from a primary DNS server or master name server to a secondary DNS server. Zone transfers ensure that there isn't a single DNS server hosting all the resource records. Windows Server 2003 supports three different types of zone transfers.

- *Full Zone Transfer (AXFR):* The entire DNS zone is replicated to the DNS server hosting the secondary DNS zone and generates more network traffic.
- *Incremental Zone Transfer (IXFR):* Only the changes to the zone file are replicated to the DNS server hosting the secondary DNS zone. This generates less traffic than a full zone transfer.
- *Fast zone transfer:* Multiple resource records can be replicated in a single message.

Zone transfers occur in the following situations:

- The DNS server is started on the secondary DNS server.
- The master name server notifies the secondary DNS servers that there are updates to the zone file.
- A secondary DNS server requests a zone transfer from its master name server.
- The refresh interval expires.

Windows Server 2003 DNS supports different types of server roles:

- Caching-only DNS server
- Does not contain zone information
- Reduces network traffic
- Primary DNS server
- Maintains the writable copy of a zone file
- Reduces network traffic
- Secondary DNS server
- Maintains a copy of a zone file
- Zone transfer from a master name Server
- Incremental zone transfers reduce network traffic
- Forwarding DNS servers
- Forwards queries that can not be resolved to specific DNS servers

See [Understanding zones and zone transfers](#) for more information.

DNS Security

When designing the DNS name resolution strategy, security must also be considered. The DNS infrastructure must be protected against various attacks. Knowing that Active Directory relies heavily on DNS, along with many other server services, you should take certain steps to secure your DNS infrastructure because there are many exploits that can take advantage of security holes. Instead of waiting for an attack to occur, you should take proactive measures to secure your DNS servers before this occurs. There are a number of ways in which a DNS infrastructure can be secured:

- Host the internal namespace on internal DNS servers. Host the external namespace on external DNS servers. If necessary, internal DNS servers can forward name resolution requests to external DNS servers.
- Place a firewall between the internal and external DNS servers.
- Store zone data in Active Directory to take advantage of secure updates.
- Use the NT File System (NTFS) instead of the File Allocation Table (FAT) file system.
- Restrict zone transfers.
- If zone data is being replicated across the Internet, data can be secured using a virtual private network (VPN) tunnel or Internet Protocol Security (IPSec).
- Use the Secure cache against pollution option shown in *Figure 5*. This prevents an attacker from polluting the zone file with invalid resource records.

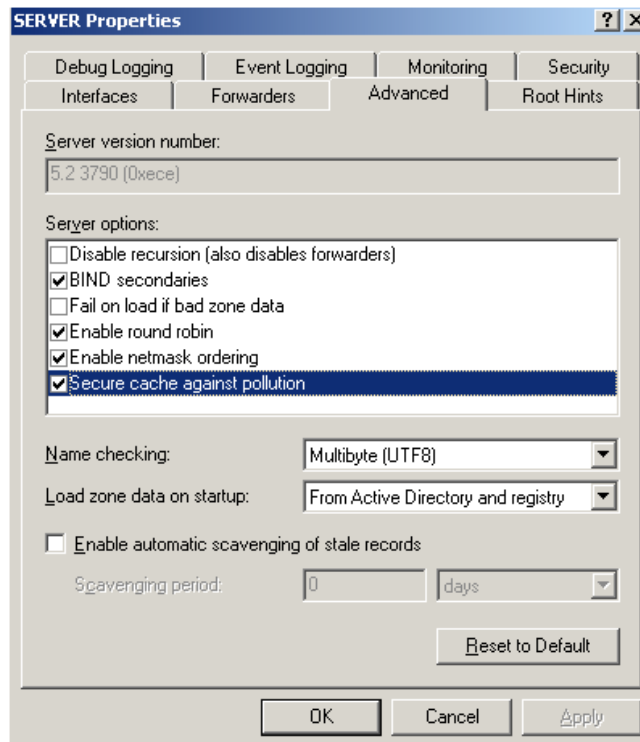


Figure 5 – The Secure Cache Against Pollution Option

See [Security information for DNS](#) for more information.

Creating the Physical DNS Design

Design DNS Server Placement

Active Directory requires DNS. Therefore, [DNS server placement](#) is an important factor in the physical design. Your overall goal should be to place DNS servers as close as possible to domain controllers and client computers on the network.

The following questions can be used to determine where DNS servers should be placed:

- How many zones are there and how large are they?
- How many DNS clients are on the network? How many DNS queries will be generated? Multiple DNS servers can provide load balancing.
- Can users resolve host names if a DNS server goes down?
- Is the network routed? A failed router can prevent DNS clients from resolving host names.

For a routed network, the ideal is to place a DNS server on each subnet. However, this is not always possible due to cost constraints. Or, if the subnets are connected via high-speed links, you can place all the DNS servers in a central location (which may be more ideal for administrators). However, in choosing the second option, the loss of a connection can mean clients are temporarily without name resolution services.

If zones are not Active Directory-integrated, consider implementing secondary DNS servers for fault tolerance. Secondary DNS servers should be placed on subnets that contain the most hosts or on those that generate the most name resolution traffic. For those subnets that generate a large amount of traffic, you may want to place both a primary and secondary server for load balancing purposes.

Also consider the speed of wide area network (WAN) connections for DNS server placement. In these situations, caching-only servers are ideal for providing name resolution to remote sites while reducing WAN traffic. Once the cache is built up, there will not be as much of a need to resolve queries over the WAN connection. The WAN link will also not be used for zone transfers (which can generate a large amount of network traffic) because the server does not store any zone information.

You also need to consider DNS replication traffic when placing DNS servers because zones must be replicated when changes occur.

Domain Controllers

Domain controllers are a crucial aspect of an Active Directory infrastructure. A domain controller is a computer running Windows Server 2003 with Active Directory installed. Each domain within a forest has its own local domain database (part of which is replicated to the global catalog server) and each domain controller within a domain maintains a working copy of the database.

A simple Active Directory implementation can consist only of a single domain controller, although multiple domain controllers can be used to provide fault tolerance and to improve response time for users. The specific functions performed by domain controllers are as follows:

- Each domain controller maintains a working copy of the local domain database.
- All domain controllers can receive updates to the local domain database and replicate the changes to other domain controllers within the same domain.
- Specific domain controllers act as bridgehead servers to replicate domain database changes to domain controllers in other sites.
- Specific domain controllers perform operation master roles.
- Domain controllers manage users' interaction on the network, such as locating and accessing Active Directory objects and validating log on attempts.

Flexible Single Master Operation (FSMO) Roles

With the multimaster replication model most changes made can be done on any domain controller and then replicated throughout the domain. These changes include day to day tasks such as adding user accounts, changing account properties, or updating policy settings. There is, however, some operations are performed with the single master replication model because they must occur in one place one only. For these types of operations, there are specific domain controllers that are designated as being operation masters and become responsible for receiving certain updates. The domain controllers that are assigned these roles then become "role owners" for the specific operations. The five different [operation master roles](#) include:

- Schema master
- Domain naming master
- Relative ID master
- Primary domain controller emulator master
- Infrastructure master

Forest Wide Operation Master

Of the five operation master roles, two of the roles are forest wide. This means that there is only one domain controller designated as the role owner for that specific operation throughout the entire forest. The two operation master roles that are forest wide are the schema master and the domain naming master.

Schema Master

Having schema updates performed on any domain controller in the forest can result in inconsistencies. The schema master is the one domain controller in the entire domain where changes to the schema can be made. It's the only domain controller in the domain where write operations can be performed to the schema and attempting to perform write operations to the schema from other domain controllers will

result in an error message. Once modifications are made, they are then replicated to other domain controllers in the forest.

The first server to have Active Directory installed automatically assumes the role of the schema master. If necessary, the role can be transferred to another domain controller using the Active Directory Schema snap-in.

Domain Naming Master

The domain naming master is primarily responsible for two tasks. It's responsible for the addition of and removal of domains to the forest. When you run the Active Directory installation wizard, the domain naming master must be available to create a new domain or remove an existing domain. If the domain naming master is unavailable, both operations will fail.

As with the schema master, the first server with Active Directory installed assumes the role of the domain naming master. The role can be transferred to another domain controller using the Active Directory Domains and Trusts snap-in.

Domainwide Operation Masters

The remaining three operation master roles are domain wide. Each domain must have one domain controller designated to perform these operations and, if needed, they can be performed on the same single domain controller. The domainwide operation master roles include the relative ID (RID) master, the primary domain controller emulator, and the infrastructure master.

Relative ID Master

The RID master is responsible for assigning strings of relative IDs to domain controllers within a domain. Each time a new object is created within Active Directory it is assigned a security ID that consists of a domain ID that identifies the domain and a relative ID that uniquely identifies the object.

When a domain controller runs out of relative IDs, it must contact the RID master before creating any new objects. Also, when you attempt to move objects between domains using the **movetree** command, the command must be initiated on the RID master in the domain where the object exists. If the RID master is unavailable, both operations will fail.

Primary Domain Controller (PDC) Emulator

The PDC Emulator acts as a Windows NT PDC for legacy clients and for any Windows NT backup domain controllers (BDCs) that may still exist on the network (chances are if you are performing a migration from Windows NT to Windows Server 2003 you will at some point have a BDC co-existing on the network). The PDC Emulator is responsible for processing password changes and replicating them to any BDCs on the network.

Once a domain has been fully migrated and the functional level has been upgraded, the PDC Emulator still receives preferential replication of any password changes made on other domain controllers in the domain. When a password is changed, it is replicated throughout the domain but can take time to propagate to all domain controllers. If a user is unable to log on due to an incorrect password, the authentication request will be forwarded to the PDC Emulator before the authentication request is denied.

Infrastructure Master

The infrastructure master is responsible for any group-to-user references when members of a group are renamed or changed. When a group contains user accounts from another domain and a user account is renamed, it takes time to propagate the changes to other domain controllers. This means that when a user account is renamed, it may take a while before the user accounts new name is displayed. The infrastructure master is responsible for updating any user to group membership changes. It makes the update locally then replicates the change to all other domain controllers within the domain. If the infrastructure master is unavailable, it takes longer for these changes to appear.

The first server to have Active Directory installed automatically assumes all three roles. To assign one of the domainwide roles to another server, use the Active Directory Users and Computers snap-in.

Although one server per forest and one server per domain assume the operations master roles, as you get in the design aspects of Active Directory, there are instances when some of the roles must be transferred to other domain controllers.

Placement of FSMO Roles

Use the following points as guidelines when planning the [placement of the flexible operations master roles](#):

- The first domain controller in the domain is designated all three roles.
- The first domain controller in the forest is designated all five roles.
- The schema master should be placed as close as possible to the Schema Admins. Only members of this group can update the schema.
- The domainwide operations masters should be located centrally. If there are multiple sites, try to avoid router hops. Place them where users and domain controllers can access them.
- PDC emulators and RID masters should be placed as close as possible to users.
- Designate specific servers as backup operations master in the event that an existing server fails.
- In a multidomain network, the Infrastructure master should never be designated as a global catalog server.

Design the Active Directory Infrastructure to Meet Business and Technical Requirements

One of the benefits of Active Directory is the ability to group objects based on the logical organization of the network, essentially masking the physical network, making it easier to locate and manage resources. Active Directory uses the following components to group resources and objects based on the logical structure of an organization:

- Forests
- Trees
- Domains
- Organizational units (OUs)

Design the Envisioned Administrative Model

- Determine the type of administrative model that is currently used within a business. (See *Table 2*.)

Administrative Model	Description
Centralized	Decision-making authority is held by a central group within the business.
Decentralized	Decision-making authority is dispersed between multiple units within the business.
Mixed	Some decision-making authority is held by a central group within the business, whereas some administrative tasks are dispersed between multiple units within the business.

Table 2 — Administrative Models and Their Descriptions

- Identify who has decision-making authority.
- Identify the type of IT structure: centralized or decentralized.
- Design a model for administration that meets business needs.
 - ▶ *Geographical:*
 - The Active Directory structure is based on geographical locations.
 - It's not affected by reorganization and expansion.
 - ▶ *Organizational:*
 - The Active Directory structure is based on business units or departments.
 - It's affected by departmental reorganizations.
 - ▶ *Functional:*
 - The Active Directory structure is based on job roles.
 - It's typically not affected by company reorganization.

- ▶ *Hybrid:*
 - The Active Directory structure is arranged geographically and then organizationally.
 - The Active Directory structure is arranged organizationally and then geographically.

See [Planning your Active Directory and Administrative Model](#) for more information.

Create the Conceptual Design of the Active Directory Forest Structure

- The forest is at the top of the Active Directory hierarchy.
- Forests basically define the boundary of an Active Directory implementation.
- In its simplest forms, a forest is a single domain.
- In more complex designs, forests can consist of a single tree with multiple domains and multiple trees.
- Active Directory forests have the following characteristics:
 - ▶ All domains within a single forest share a common schema.
 - ▶ All domains within a single forest share a common global catalog.
 - ▶ Two-way transitive trust relationships exist between parent and child domains and domain trees.
 - ▶ There is a single Schema Admins group and Enterprise Admins group for the entire forest.
 - Members of the Enterprise Admins group have administrative privileges and rights to all domains within the forest.
- Create multiple forests if:
 - ▶ You must maintain limited trusts.
 - ▶ Multiple global directories are required.
 - ▶ Multiple schemas are required.
- A tree can be defined as a hierarchical arrangement of domains within a forest.
 - ▶ A tree is established when a child domain is added to the hierarchy under an existing parent domain.
 - ▶ A single forest can contain multiple trees.
 - ▶ When you create a new tree within a forest, you are establishing a new namespace.
 - ▶ One of the main reasons for creating multiple trees is to maintain separate namespaces while still maintaining a common schema and global catalog, as shown in *Figure 6*.

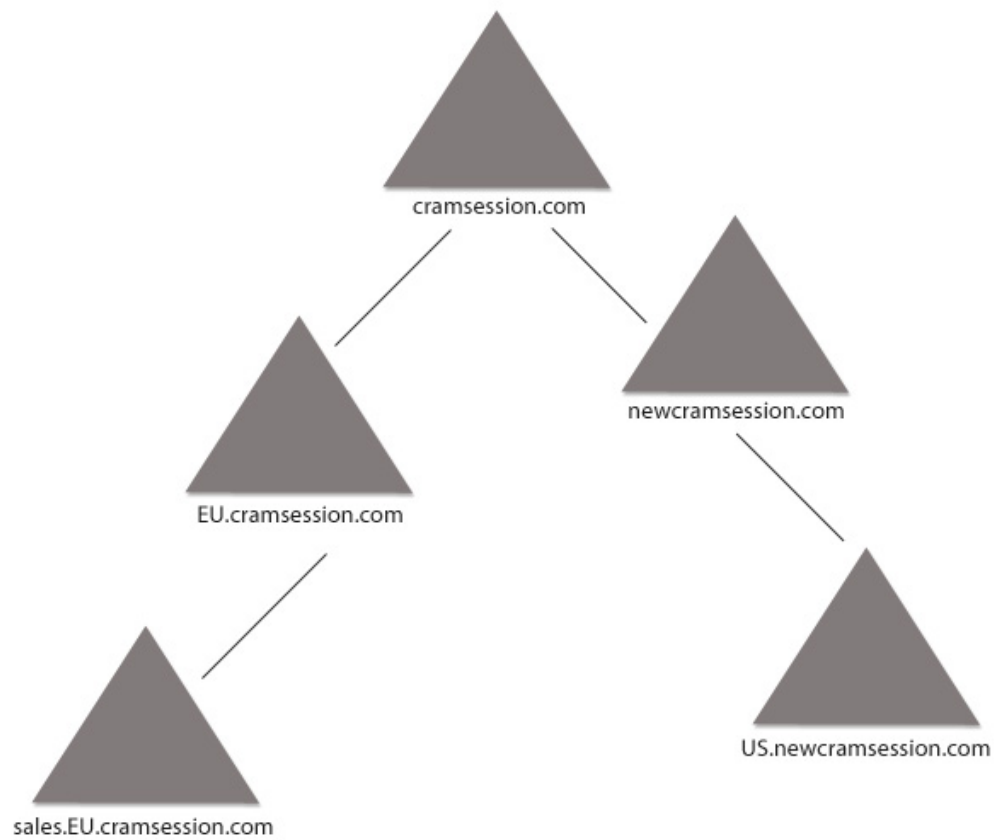


Figure 6 — Multiple Trees

Create the Conceptual Design of the Active Directory Domain Structure

- Domains are the core unit of an Active Directory structure.
 - ▶ Domains establish security boundaries within a forest.
 - They ensure that administrators have only the rights to perform administrative tasks and manage resources within their own domain.
 - ▶ Domains as a unit of replication:
 - Each domain within a forest has its own local domain database.
 - All domain controllers within the domain maintain a full working replica of the database.
 - Any changes made to the database are replicated to all other domain controllers within the domain.
- The first domain is the forest root domain.

- Create multiple domains:
 - ▶ To maintain distinct administrative boundaries
 - ▶ Decentralized administration (OUs can also be used)
 - ▶ Multiple security policies
 - ▶ Unique namespaces
 - ▶ Optimize network traffic

Trust Relationships

When it comes time to design an Active Directory structure, you must understand how trusts relationships are implemented. Trust relationships are important because they provide users with the means to access resources in other domains.

There are basically four different types of [trust relationships](#) that exist in a Windows Server 2003 environment: two-way transitive trusts, shortcut trusts, forest trust, and external trusts.

- *Two-way transitive trusts:*
 - ▶ When a new domain is added to the forest, a two-way transitive trust is automatically created between the new domain and the parent domain.
 - ▶ When a new root domain is added to an existing forest establishing a new tree, a two-way transitive trust is configured between the forest root and the root domain of the new tree.
- *Shortcut trusts:*
 - ▶ Created to decrease authentication time between domains in the same forest.
 - ▶ As shown in *Figure 7*, a shortcut trust is a two-way transitive trust between two domains only, it must be explicitly defined by an administrator.

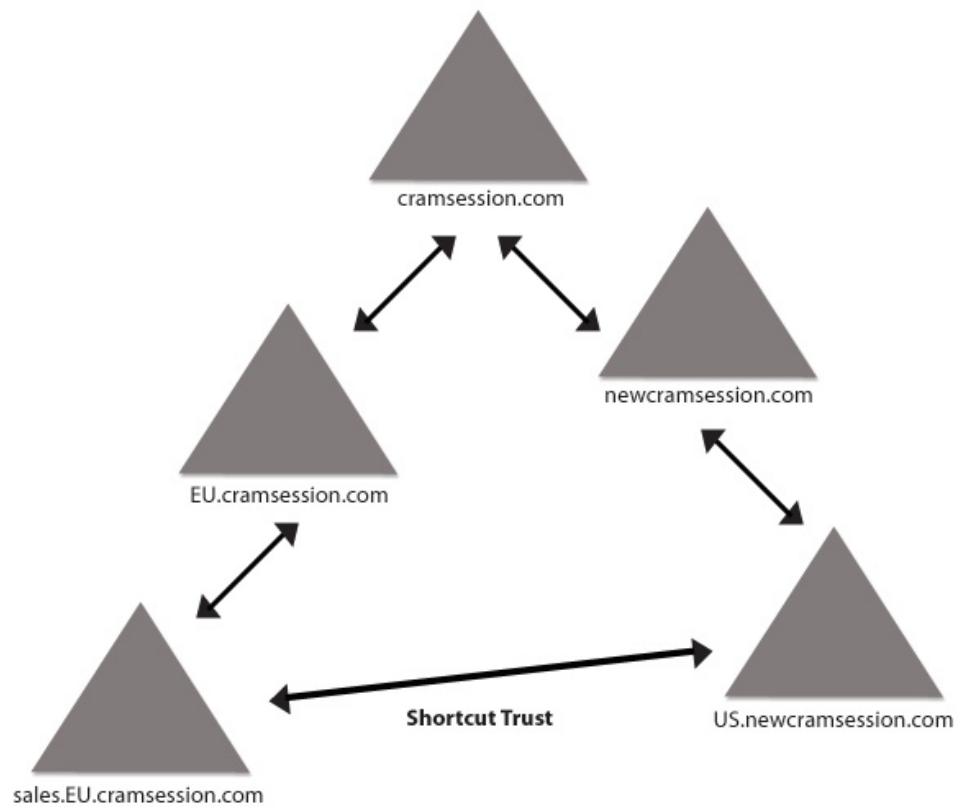


Figure 7 - Shortcut Trust

- *Forest trusts:*
 - ▶ Extends the two-way transitive trusts beyond a single forest to a second forest.
 - ▶ Links two separate forests together, creating two-way transitive trusts between domains.
 - ▶ When linking more than two forests together, keep in mind that a forest trust is not transitive.
 - ▶ Functional level of the forest must be raised to Windows Server 2003 mode to use forest trusts.
- *External trusts:*
 - ▶ Provide users in a Windows Server 2003 domain access to resources in a Windows NT 4.0 domain.
 - ▶ If two forests are not joined in a forest trust relationship, create an external non-transitive trust between the source domain and the destination domain.

See [Trust Technologies](#) for more information.

Create the Conceptual Design of the Organization Unit (OU) Structure

- OUs are containers objects used to logically organize objects within a domain for administrative purposes.
 - OUs can contain objects, such as printers, computers, user accounts, shares, or other OUs.
 - Each domain within a forest can implement its own OU hierarchy that is completely independent of all others, as shown in *Figure 8*.

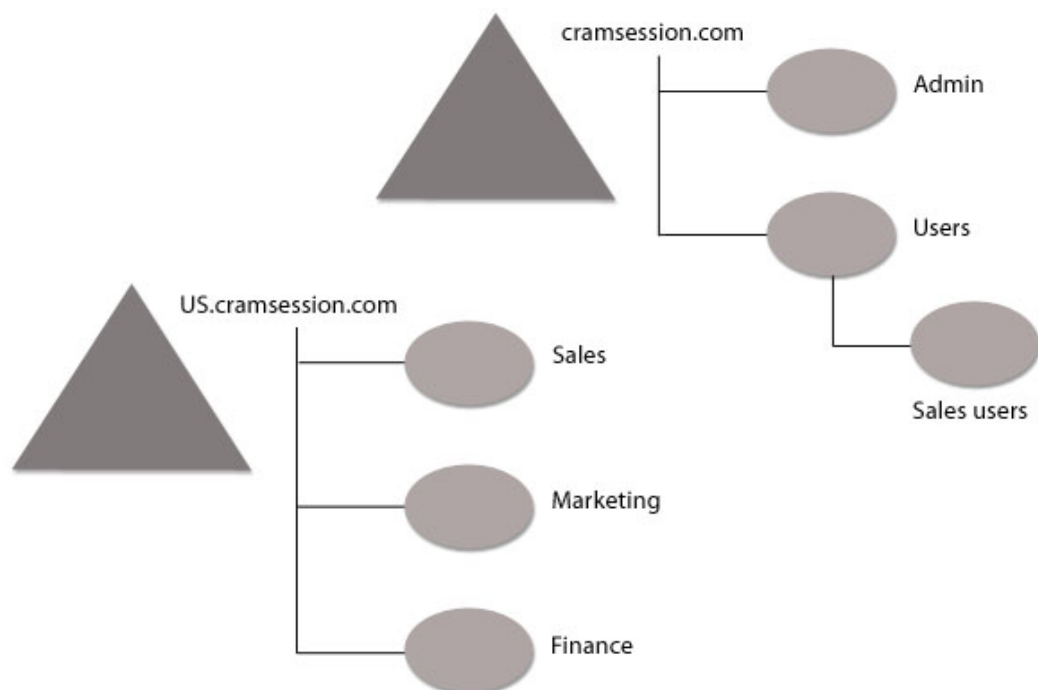


Figure 8 - OU Structure

- One of the main reasons for creating OUs to group objects is for delegation of control.
- Upper layers of the OU structure should be based on the model for administration.
- Lower level OUs should be created for administrative purposes.
 - Administrators must be able to apply GPOs to the necessary objects.

See [Reviewing Organizational Unit Design Concepts](#) for more information.

Design the Network Services Infrastructure to Meet Business and Technical Requirements

Create the Conceptual Design of the DNS Infrastructure

- Identify the DNS namespace that will be used.
 - ▶ External names should be registered on the Internet.
- Identify where on the network the DNS servers should be placed.
 - ▶ In a routed network, DNS servers should be placed on all subnets or on the subnets with the most number of DNS clients.
 - ▶ Secondary servers should be implemented for fault tolerance.
 - ▶ Caching only servers are ideal for remote sites.
 - ▶ If remote sites contain a large number of users, consider using a full DNS implementation.
- Identify any interoperability issues with older versions of DNS or BIND.

See [How DNS Support for Active Directory Works](#) for more information.

Create the Conceptual Design of the WINS Infrastructure

- WINS is not required if all clients and applications support DNS.
- WINS resolves NetBIOS names to IP addresses.
- On small subnets with few clients, WINS may not be required.
 - ▶ Can use LMHOSTS or broadcasts
- Using broadcasts for name resolution can have an impact on network performance.
- WINS reduces the number of broadcasts on a subnet.
- Identify the number of WINS servers required.
 - ▶ Number of NetBIOS registrations
 - ▶ Frequency of registrations, renewals, and resolution requests
 - ▶ Physical network connections
 - ▶ Recommended that there are two WINS servers for every 10,000 clients

Create the Conceptual Design of the DHCP Infrastructure

- DHCP is used to dynamically assign IP addresses to computers.
- Number of DHCP servers required:
 - ▶ Single server creates a point of failure
 - ▶ 80/20 rule for fault tolerance between DHCP servers
 - For each range of IP addresses, one DHCP server is configured with 80% of the IP addresses whereas a second DHCP server is configured with the remaining 20%

- ▶ Automatic Private IP Addressing (APIPA) if no DHCP server is available
 - Cannot communicate outside of the local subnet
- ▶ DHCP relay agent to extend DHCP functionality across subnets
- ▶ Physical network connections
 - Use multiple DHCP servers for slow WAN connections
- Identify where DHCP servers will be placed:
 - ▶ Place DHCP servers on the subnet with the most DHCP clients.
 - ▶ Place a DHCP server on either side of a slow WAN connection.
- Routed network:
 - ▶ DHCP relay agents pass DHCP messages between clients and DHCP servers located on different subnets.
 - Required if routers are non-RFC 1542 compliant
 - ▶ RFC 1542-compliant routers can forward DHCP broadcasts.
- All DHCP servers can be placed in a central location (centralized):
 - ▶ Administration of DHCP might be simpler because there will be fewer servers. Having them in a single location where technical support is readily available makes problems easier to troubleshoot.
 - ▶ If fewer DHCP servers are required, the cost associated with implementing DHCP will be reduced.
 - ▶ Clients in remote sites might have to rely on WAN links to obtain an IP address. That means if the WAN link is down, clients will not be able to contact a DHCP server. This creates a single point of failure.
 - ▶ Administration could be more difficult because it is hard for an administrator to know the configuration requirements of a remote site.
- DHCP servers can be placed on different subnets (decentralized):
 - ▶ Local administrators can administer their own DHCP servers and configure them in a way that meets their own needs.
 - ▶ Clients do not have to rely on WAN links to obtain an IP address.
 - ▶ Placing a DHCP server at each location results in an increase in cost because multiple servers are required.

See [Planning DHCP networks](#) for more information.

Create the Conceptual Design of the Remote Access Infrastructure

- There are two remote access methods: dial-up and VPN.
- Windows Server 2003 supports two line protocols:
 - ▶ Point-to-Point Protocol (PPP)
 - ▶ Serial Line Internet Protocol (SLIP)
- VPN creates a secure tunnel over a public network such as the Internet

- Windows Server 2003 supports two tunneling protocols:
 - Point-to-Point Tunneling Protocol (PPTP)
 - Layer Two Tunneling Protocol over Internet Protocol Security (L2TP/IPSec)

Design an OU Structure

OUs are created to logically organize objects within a domain. They are also created for the purpose of delegating authority and applying group policies.

Identify the Group Policy Requirements for the OU Structure

- Determine what areas require Group Policy.
- This impacts the design of lower level OUs.
- Identify what areas of the users' computing environment must be managed.
- Identify what levels of administration are required throughout a business. Some areas may require high levels of management whereas others may only require a minimal level.

Design an OU Structure for the Purpose of Delegation

- [Delegation](#) is the process of assigning a user or group the ability to perform specific administrative tasks.
- Determine who will be assigned privileges and the scope of the privileges.
- Delegate control at the OU level to limit the scope of the assigned privileges:
 - Privileges delegated on a parent OU are inherited by any child OUs.
- The OU structure should allow a business to easily delegate control.

Design a Security Group Strategy

- Distribution groups are used to send e-mail messages to multiple users at a time.
- Security groups have security descriptors associated with them and are used to control access to resources or apply user rights.
- There are four different types of [security groups](#) that can be created: local groups, domain local groups, global groups, and universal groups.
- The different groups can be characterized by their scope.
- The scope of a group determines how it can be used in the forest and the members it can contain.

Define the Scope of a Security Group to Meet Requirements

- Local groups:
 - Local groups are created to assign rights and permissions to a local computer only.
 - A local group can contain local user accounts, domain user accounts, computer accounts, and global groups.
 - Once you install Active Directory and promote the server to a domain controller, you must use the Active Directory Users and Computers snap-in to create either global, domain local or universal groups.
- Global groups:
 - Global groups logically organize user and computer accounts, sometimes based on departments or business functions.
 - A global group can contain members only from the domain in which it was created.
- Domain local groups:
 - Domain local groups are used to grant permissions to resources within a single domain.
 - In terms of scope, they can contain accounts and groups from any domain in any forest.
 - Domain local groups can be used only to assign permissions to resources within the domain that the group is created.
- Universal groups:
 - Universal groups are used to combine groups from multiple domains in a single forest.
 - Universal groups can contain accounts and groups from any domain in any forest and can be assigned permission to any resources within the forest.
 - This group has the greatest scope.

Changing Group Scopes

When changing the scope of a group, there are certain limitations and restrictions:

- A global group can be converted to a universal group as long as it has not been nested within another global group.
- A universal group can be converted to a global group as long as the group to be changed has not been nested within another universal group.
- A domain local group can be converted to a universal group as long as the group to be changed does not contain any other domain local groups as members.
- A universal group can be converted to a domain local group with no restrictions.

Nesting

- [Group nesting](#) allows you to add groups within groups. In other words, groups can be members of other groups.
- The only thing to keep in mind is to use it wisely and limit the level of nesting as it can soon get out of hand and difficult to track.
- The nesting options are dependent on the mode in which your domain is running.

- ▶ If your domain is running in mixed mode, group nesting is available, but there are limitations.
- In terms of group nesting, if the domain is running in Windows 2000 native mode:
 - ▶ Universal groups can contain other universal groups or global groups from any domain in the forest.
 - ▶ Global groups can contain other global groups from within the same domain.
 - ▶ Domain local groups can contain global groups from any domain in the forest, universal groups, and other domain local groups from within the same domain.
- If the domain is running in mixed mode, the following restrictions apply:
 - ▶ Global groups cannot have other groups as members. They can contain only user accounts.
 - ▶ Domain local groups can have only global groups and user accounts as members.

Define Resources Access Requirements

The recommended strategy for implementing groups is as follows:

1. Start by organizing users with common needs or common job functions into global groups.
2. Create domain local groups in each domain to grant access to shared resources.
3. Add the global groups that require access to shared resources to the appropriate domain local groups.
4. Assign the necessary permission to the domain local group.

Define Administrative Access Requirements

- Use the principle of least privilege when designing administrative access requirements.
- Use delegation of authority instead of adding user accounts to the Administrators group.
- Assign the minimum privileges required to perform any administrative tasks.

Define User Roles

- Define user roles for access control.
- Define groups based on user roles for assigning permissions.
- It's easier to track permissions when they have been assigned to groups instead of users.

Design an Active Directory Site Topology

The physical structure of a network is defined using sites. Active Directory sites are created to optimize replication between domain controllers that are connected by slow, unreliable, and heavily utilized links.

- A site can be defined as a group of IP subnets that have fast, reliable connectivity.
- Group IP subnets with fast, reliable connectivity into a single site.
- Use separate sites when IP subnets have unreliable, slow, or heavily used connections.
- Replication between sites can be controlled.

Design Sites

- Windows NT used a single master replication model where changes had to be made on the PDC and then replicated to the BDCs.
 - ▶ There was real no way to control replication within a domain.
 - ▶ Change the pulse setting in the Windows NT registry.
- Windows Server 2003, like Windows 2000, supports [multimaster replication](#).
 - ▶ There are no PDCs or BDCs.
 - ▶ All Windows Server 2003 domain controllers are peers and any changes made on a domain controller are replicated to all others in the domain.
- Sites can optimize network traffic and control how/when replication occurs between different domain controllers that belong to the same domain but in different sites.
- Sites can be created to control the following types of traffic:
 - ▶ *Replication traffic* - Sites can be created to control how changes are replicated between domain controllers in another site.
 - ▶ *Network logon traffic* - When a user logs onto the network, the user is authenticated by a domain controller in the same site, eliminating the need for a user to log on over a slow connection.
 - ▶ *Distributed file system (Dfs)* - If DFS is implemented on the network, when a user attempts to access a share within the DFS hierarchy, they will be directed to a server within their own site.
 - ▶ *File Replication Service* - The FRS is site-aware and uses the site topology when replicating information such as the SYSVOL, group policy settings and user logon/logoff scripts.

When you place a group of IP subnets into a single site, it is assumed that they have good connectivity and can handle regular traffic generated from Active Directory replication, as opposed to placing IP subnets into different sites. Therefore, replication within a site occurs differently than it does between sites.

Intra-Site Replication

- Replication within a site takes place more often than it does between sites.
- Within a site, replication occurs automatically.
- When you install Active Directory a default site is created.
- A replication path is automatically generated by the Knowledge Consistency Checker service forming a ring topology between domain controllers.
- The replication topology defines how updates to the local domain database flow between domain controllers.
- The topology that is generated ensures that there are always two replication paths between domain controllers thereby ensuring replication can still occur if a domain controller is unavailable, as shown in *Figure 9*.




Figure 9 - Intra-site Replication

- Within a site, domain controllers notify each other when changes to the local domain database occur.
- Domain controllers poll each other at regular intervals to check for changes and updates.
- Replication can occur between any domain controllers.
- [Intra-site replication](#) uses remote procedure calls (RPCs) to replicate information.

Intersite Replication

- Sites and site links define how Active Directory replicates information between domain controllers on different subnets.
- [Intersite replication](#) has the following characteristics:
 - ▶ Domain controllers in different sites do not notify each other when changes occur.
 - ▶ Domain controllers poll one another at preconfigured intervals during specific times to check for updates and changes.
 - ▶ Replication occurs between specific domain controllers (known as bridgehead servers), then replicated to other domain controllers within the site.
 - ▶ Information replicated across site links can be compressed.
 - ▶ Replication can use Remote Procedure Call (RPC) over Transmission Control Protocol/Internet Protocol (TCP/IP) or Simple Mail Transfer Protocol (SMTP).

Identify Site Links

- Network connections between sites are represented within Active Directory as [site links](#).
- Site links are the logical connections between sites and must be created before replication can occur.
- Once a site link is defined, replication occurs based on the attributes of the link.
- Each site link will have to following attributes defining when and how replication can occur across the link:
 - ▶ Replication interval
 - ▶ Replication schedule
 - ▶ Replication transport
 - ▶ Link cost
- Site links are transitive in nature.

Replication Interval

- There is no notification process that occurs between domain controllers in different sites when changes occur.
- Specific domain controllers in each site poll each other for changes.
 - ▶ These servers are called Bridgehead Servers.
- The replication interval defines how often a domain controller in one site can use the connection to poll a domain controller in another site for changes.
 - ▶ By default, the replication interval value is set to 15 minutes.

Replication Schedule

- The replication schedule for a site link defines when the connection can be used for replication.
- The replication schedule tells Active Directory when the connection can be used for replication and the interval defines how often during the period of availability, domain controllers' poll for updates.

Replication Transport

- There are two methods available to transmit information across a site link: using RPC over TCP/IP or using the SMTP.
- RPC:
 - ▶ The default transport for inter-site replication (as well as intra-site replication).
 - ▶ RPCs are synchronous in nature.
 - ▶ Requires a direct connection with the destination server before replication can occur.
- SMTP:
 - ▶ Does not require a direct connection and replicates information as simple mail messages.
 - ▶ It uses a store-and-forward method of replication information.
 - ▶ The advantage of this is if the destination server is unavailable, the information to be replicated can be stored and forwarded when the server is available.
 - ▶ SMTP does not replicate the domain partition.

Link Cost

- [Link costs](#) allow you to control how information flows between sites or the path information travels.
- A link cost is a numeric value assigned to a site link. If there are multiple connections between sites, multiple site links can be created, and a cost assigned to each one that reflects the speed and available bandwidth.
- When replicating information, Active Directory will always choose the link with the lowest cost, unless it is unavailable.

New Features in Windows Server 2003

Active Directory

Active Directory is the core feature of Windows Server 2003 providing a central location to store information about objects within a network making it easier for administrators to manage and makes it simpler for users to locate information. Some of the new features and improvements introduced in Windows Server 2003 Active Directory include:

- *Cached credentials* – In the event that a Global Catalog server is unavailable users can still log on to a local domain controller using cached credentials. This is great for those branch offices and remote locations connected to a global catalog server via a slow or unreliable connection.
- *Improved replication and synchronization* – With Windows Server 2003 administrators now have more control over the information that is replicated and synchronized between domain controllers, both within a domain and between domains. Windows Server 2003 also includes WMI providers for monitoring and alerting administrators to problems with intra-domain replication and inter-domain trusts.
- *Saved queries* – Administrators can now save Active Directory queries and export them in XML for future reference.
- *Cross forest and trust management* – Users can now enjoy the benefits of a single log on when accessing resources in other forests.
- *Schema* – Classes and attributes added to the schema can now be deactivated. This is useful if the schema administrator makes an error when adding information to the schema. The attribute or class that was incorrectly entered can be deactivated and the correct information entered.
- *Passport integration* – IIS 6.0 now supports passport integration mapping objects stored in Active Directory to their corresponding passports.

Security

Nowadays one of the hottest topics is security especially with more and more organizations connecting to the Internet using it as a business tool exposing their internal network to the possibility of viruses and hackers. Administrators are always looking for new security features to implement on the network. Whether it's standardizing the users computing environment or protecting the network from external intrusion, new security features and tools are always welcome. Windows Server 2003 comes with many new and improved security features that are effective and simple for administrators to implement. Some of these new features and improved features include:

- *Effective permissions* – For those of you who have at one time or another had to troubleshoot a permission problem, you know how difficult it can be, depending on the number of users and the number of shares. If a user is having access problems the task of checking the permissions on the share as well as all the parent containers for the user account begins and then doing the same thing for any groups the user is a member of. Windows Server 2003 introduces the effective permissions tool that will calculate a user's or groups effective permissions taking into account both permission inheritance and group membership. The only important thing to keep in mind with this tool is the calculation does not take into account any share permissions.

- *Software restriction policies* – This is a great new feature that allows an administrator more control over the users' computing environment. It allows you to specify the type of software users are permitted to run, with the intent being to protect them from running damaging software such as a Trojan horse.
- *Stored user names and passwords* – This feature allows users to connect to servers using credentials other than those they are currently logged on with.

Analyze the Impact of Active Directory on the Existing Technical Environment

Analyze Hardware and Software Requirements

Analyze the existing hardware and software that is being used to help in determining the impact that Active Directory will have on the existing technical environment.

- Determine if the existing hardware being used.
 - This will help to determine if the existing hardware meets the hardware requirements of Windows Server 2003.
- Identify the operating systems and service packs that are installed on computers.
 - Some operating systems may not be able to interoperate with Windows Server 2003 and Active Directory
 - Identify those operating systems that have to be upgraded.
- Identify the applications that the business currently uses.
 - If necessary, these applications must be tested to ensure that they will function under Windows Server 2003.

Analyze Interoperability Requirements

Businesses sometimes use different operating systems. This often occurs when a business is in the process of upgrading existing systems. Identify the current operating systems being used and how they will interoperate with Windows Server 2003 and Active Directory.

- Windows Server 2003 can support and interoperate with a variety of different operating systems: Windows 2000, Windows XP, Windows Me, Windows 9x, Windows NT 4.0, Windows 3.x, UNIX, and Macintosh.
- Windows 95 clients require the Active Directory client to interoperate with Windows Server 2003 Active Directory.
- Windows NT 4.0 clients require Service Pack 4.
- The Active Directory client can be installed on Windows NT 4.0 if Service Pack 6a has been added.
- Any workstations running Windows for Workgroups must be upgraded.

Analyze Current Level of Service within an Existing Technical Environment

The existing technical environment must be analyzed when planning for Active Directory. Some businesses rely on internal staff to provide technical support, whereas others outsource some or all of their technical support to external companies. Identify whether the existing support staff has the skills and knowledge required to support the new infrastructure. Internal support staff may require additional training.

Also consider how end users will be impacted by the new infrastructure. Some basic training as to some of the changes end users will see when Active Directory is in place should be considered.

Analyze Existing Network Operating System Implementation

During this portion of the analysis phase, identify the existing domain model that is currently employed. This will have a major impact on the Active Directory design. Also identify the how many domain controllers currently exist and where on the network they have been placed. The business will more than likely have various other servers on the network. They must be identified and their configuration documented.

Identify the Existing Domain Model

- Windows NT 4.0 domain can store up to 40,000 objects. Active Directory has no limit.
- Windows NT 4.0 supports four different domain models. (See *Table 3*.)
 - Single domain
 - Single master
 - Multiple master
 - Complete trust

Domain Model	Description
Single domain	A single domain that contains both the user accounts and resources.
Single master	User accounts are contained within a single domain as shown in <i>Figure 10</i> . Resources are contained within a separate domain. Requires trust relationships between master domain and resource domains.
Multiple master	Two or more master domains that contain user accounts as shown in <i>Figure 11</i> . Resources are contained within a separate domain. Requires two-way trust relationships between master domains. One-way trust relationships are established between master domains and resource domains.
Complete trust	Multiple domains hosting both user accounts and resources. Two-way trusts are required between each domain. This model is the most complex.

Table 3 - Domain Models and Descriptions

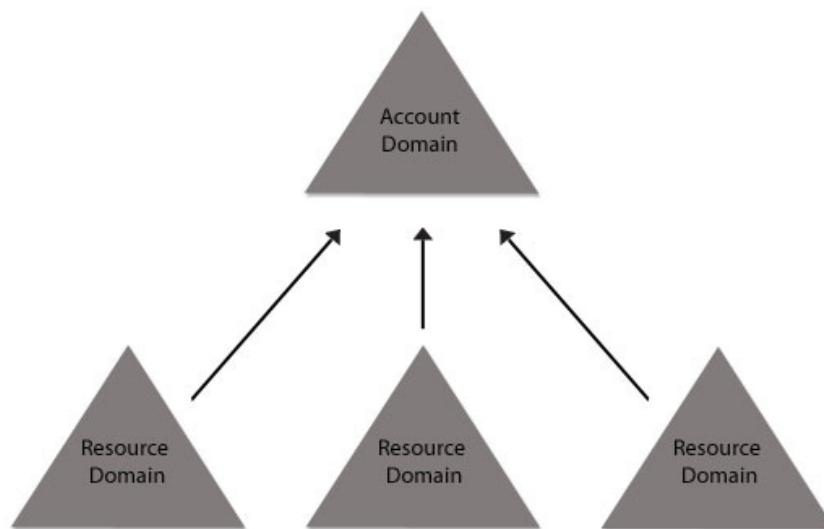


Figure 10 - Single Master Domain Model




Figure 11 - Multiple Master Domain Model

- Resource domains can be upgraded to organizational units.
- Master domains can be combined into a single Windows Server 2003 Active Directory domain.

Identify the Number and Location of Domain Controllers on the Network

[Domain controllers](#) are required for the logon process. Therefore, their availability is crucial. During the analysis of the existing environment, identify the number of domain controllers on the network and where they are placed. This information will help to determine if additional domain controllers are required to support the new infrastructure.

Identify the Configuration Details of all Servers on the Network

- Document configuration of existing servers on the network.
- Identify how the configuration will be affected by upgrading to Windows Server 2003.

Analyze Security Requirements for the Active Directory Service

The existing level of security must be maintained, or even increased, with the new infrastructure. Therefore, the security requirements of the business must be identified.

Security requirements will impact on the logical design of Active Directory.

Analyze Current Security Policies, Standards, and Procedures

- Identify configuration standards
 - ▶ Desktop configuration for end users
 - ▶ Logon restrictions
 - ▶ Application restrictions
 - ▶ Identify which users and groups require access to data
 - ▶ Public key infrastructure (PKI) infrastructure implementation
 - ▶ Remote connection authentication
 - ▶ Data encryption
- Identify account policies
 - ▶ Password policies
 - Password policies are configured at the domain level
 - Unique password policies within a business will result in a multiple domain infrastructure
 - ▶ Audit policies
 - ▶ Account lockout policies
- Identify security procedures employed by the business

Identify the Impact of Active Directory on the Current Security Infrastructure

Once the existing security has been identified, you can analyze the impact that Active Directory will have. Some security aspects may not change with Active Directory. Others, however, security aspects may change once Active Directory is implemented.

Identify the Existing Trust Relationships

- Trust relationships define the domain boundaries that users can cross to access network resources.
- Windows NT 4.0 only supports one type of trust relationship—one-way non-transitive trusts—as shown in *Figure 12*.
- Trusts are not created automatically.

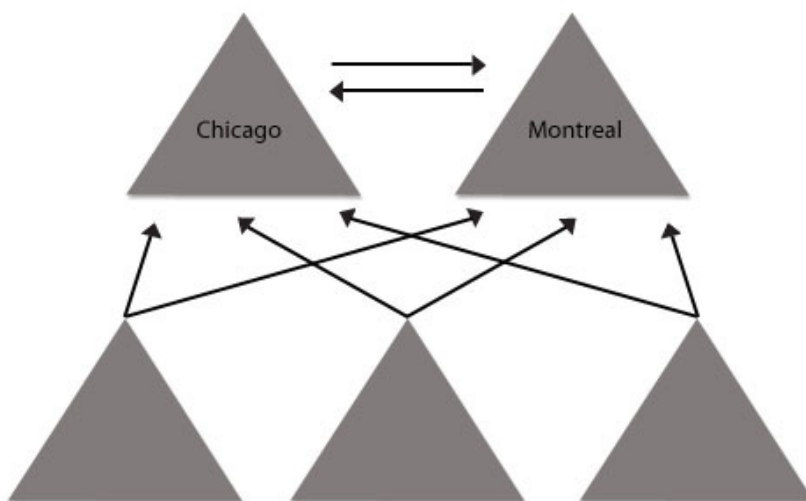


Figure 12 - One-way Non-transitive Trusts

- Windows Server 2003 trusts:
 - There are basically four different types of trust relationships that exist in a Windows Server 2003 environment including two-way transitive trusts, shortcut trusts, forest trust, and external trusts.
 - Within a single forest, administrators do not need to establish trust relationships with domains in the same forest. When a new domain is added to the forest a two-way transitive trust is automatically created between the new domain and the parent domain. As well, when a new root domain is added to an existing forest establishing a new tree, a two-way transitive trust is configured between the forest root and the root domain of the new tree. (See *Figure 13*.)

- ▶ Essentially what this does is create a trust path throughout the forest allowing users to access resources in other domains. Keep in mind as well that when a user attempts to access resources in another domain, the trust path must be followed (unless a shortcut trust has been defined between the source and destination domains).

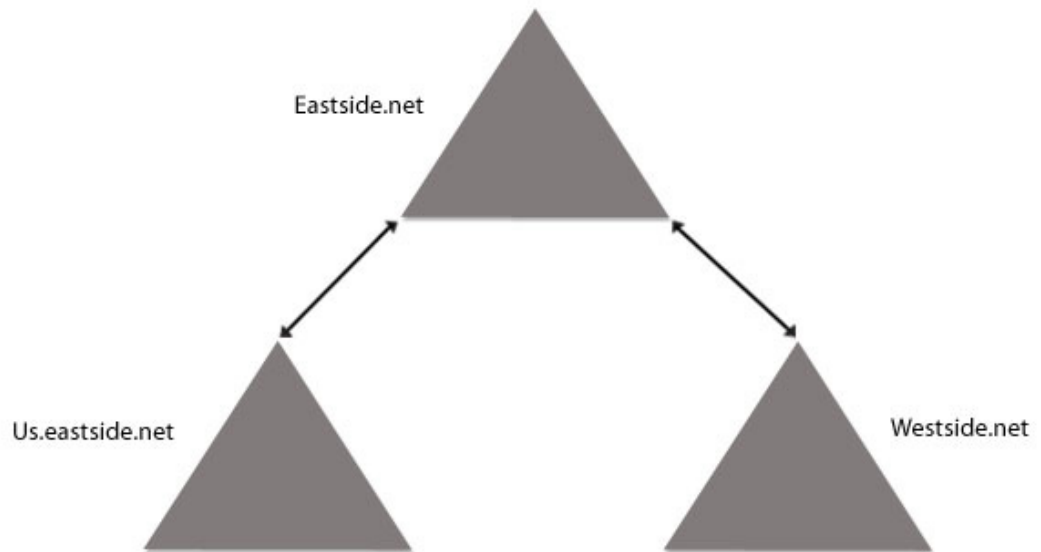


Figure 13 - Two-way Transitive Trusts

- ▶ Shortcut trusts are created to improve user response time. As already mentioned, the trust path must be followed when accessing resources outside the local domain. This means a user may have to go through several domains before reaching the domain in which the resource is located. To increase response time for authentication, you can create a shortcut trust between the source domain and the destination domain. A shortcut trust is a two-way transitive trust between two domains only it must be explicitly defined by an administrator.
- ▶ In Windows 2000, in order to provide users in one forest access to resources in another forest, a one-way trust had to be configured between the source domain and the destination domain. All this did was grant users in the source domain access to resources in the destination domain, not vice versa. It also did not give users access to any other domains other than the domain specific in the trust relationship. In Windows Server 2003 it is now easier to link two forests together as you can create a trust between two forests known as a forest trust.
- ▶ The purpose of a forest trust is to extend the two-way transitive trusts beyond a single forest to a second forest. Implementing a forest trust links two separate forests together, creating two-way transitive trusts between domains. When linking more than two forests together keep in mind that a forest trust is not transitive. This means if a forest trust is created between A and B and a forest trust is created between B and C, there is no forest trust relationship between A and C.

- ▶ The fourth type of trust in Windows Server 2003 is an external non-transitive trust. This type of trust is created for two purposes. First, it can be created to provide users in a Windows Server 2003 domain access to resources in a Windows NT 4.0 domain. Or, if two forests are not joined in a forest trust relationship, you can create an external non-transitive trust between the source domain and the destination domain.

Identify Network Topology and Performance Levels

The existing physical aspect of the network also needs to be assessed. This will help to determine if any upgrades are required before implementing the new network infrastructure.

Analyze the performance of the physical network connections and the performance of servers on the network.

Identify Constraints in the Current Network Infrastructure

- LAN/WAN connectivity
 - ▶ Identify and document the physical topology of the network.
 - ▶ Include connectivity type, speed of each connection, and bandwidth usage.
- Network Monitor can be used to determine the amount of network traffic.
 - ▶ Network Monitor is a miniature network analyzer, or sniffer.
 - ▶ A *sniffer* is a tool that can be plugged into a network that will track every packet that hits the network cable at that point.
 - ▶ The sniffer can decipher the packets, determine where they come from and where they're going, and also the contents of the packets themselves.

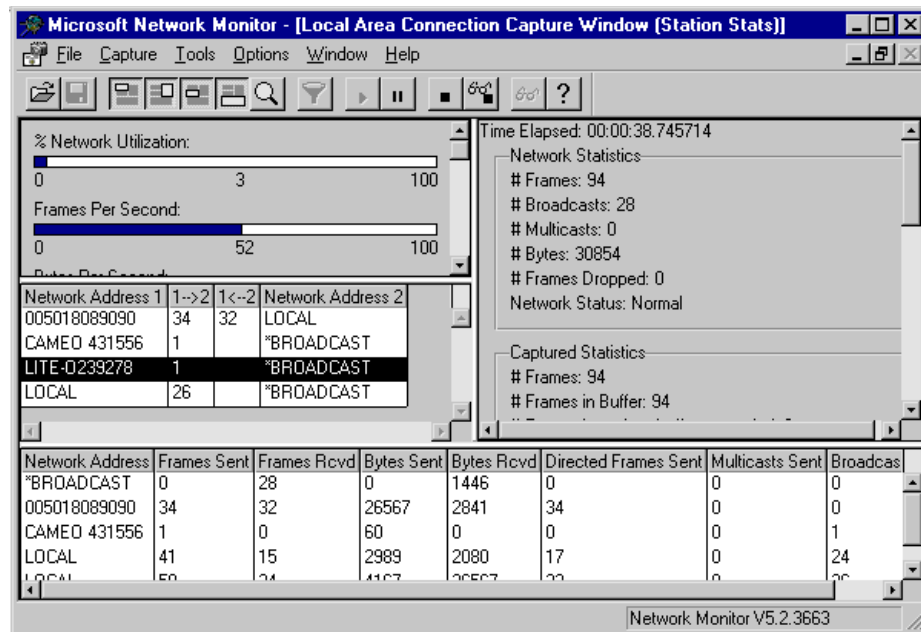


Figure 14 - Network Monitor

Interpret Current Baseline Performance Requirements for Each Major Subsystem

- A baseline is simply a set of data that depicts the "norm" for a particular object, event, or status.
- [Performance Monitor](#) works by narrowing the focus onto a specific *object*—usually a piece of hardware or major process—and then further focusing upon a specific activity associated with that object, called a counter.
 - ▶ Counters are the actual numbers, data, and statistics that are used to determine what's going on with the object.
- Include the major subsystems when establishing a baseline of performance.
 - ▶ Disk monitoring is divided into two sections: physical disk and logical disk.
 - Information about a logical disk (volume or partition) may indeed be more important than the physical disk, because it can point to issues, corruptions, and performance issues that may be able to better allow you to isolate an offending application or piece of hardware.
 - ▶ Memory is RAM.
 - Memory is used in several different ways, and counters exist for this object that can tell you how efficiently your RAM is working for you.
 - ▶ Network interface
 - This particular object will give you statistics on how many network packets are being sent and received by the NIC. It will also give you statistics about bad packets and data that are being sent and received by the same NIC.

- ▶ Paging file
 - Paging is a technology that has been used for many years as an add-on or crutch for systems with insufficient RAM. This is done by literally shifting data from RAM and storing it on the hard disk. The 'conscious' system still references the stored RAM as if it were really located on the physical RAM, but in reality, the system's 'subconscious' takes care of moving and addressing data from RAM to storage, and back. More recent versions of Windows require a paging file, even when sufficient memory exists. The purpose mainly lies in some obscure capability to increase speed and enhance productivity. As a result, the paging file and processes that surround it are thrust into the limelight of performance concerns, because an ailing paging file can lead to performance issues just as much as a healthy one can enhance performance.
- ▶ PhysicalDisk
 - The physical hard disk is another object that should be monitored occasionally for obvious reasons: a bad hard disk will generate errors. An overloaded hard disk will also have problems and could very well slow down the overall server productivity—especially if the paging file is being heavily used.
- ▶ Processor
 - The processor object has the ability to monitor the health and welfare of a server's one or more CPUs. Although a CPU may seem to be an "either it works or it doesn't" item, it's not. As with any of these other objects, counters exist that will monitor things such as how many calculations are being performed, how many errors are being generated, and even how long the line of requests is for processes waiting to for CPU attention. (Any waiting line in Performance monitor is termed a queue.)

Analyze the Impact of the Infrastructure Design on the Existing Technical Environment

Analyze Hardware and Software Requirements

- Windows Server 2003 Standard Server hardware requirements:
 - ▶ 133 MHz processor (550 MHz recommended).
 - ▶ 128 MB of RAM (256 MB recommended).
 - ▶ 1.5 GB of free disk space.
- Document information about the existing hardware.
 - ▶ Hardware must meet the minimum hardware requirements.
 - ▶ Hardware must be supported by Windows Server 2003.
 - ▶ Identify those operating systems that must be upgraded.
- Document operating systems and applications currently being used.
 - ▶ Operating systems.
 - ▶ Applications.

Analyze Interoperability Requirements

Interoperability issues can arise with hardware, operating systems, and applications. Once the existing software and hardware has been identified, the next step is to determine how they will interoperate with Windows Server 2003 and Active Directory.

Interoperability issues should be identified before rolling out the new infrastructure.

- Identify whether the existing operating systems will interoperate with Windows Server 2003 and if any additional software is required to do so.
- Identify if applications can interoperate with Windows Server 2003 and Active Directory.
- Verify that hardware is on the Hardware Compatibility List (HCL) or in the Windows Catalog.

Analyze Current Level of Service within the Existing Technical Environment

Identifying the current level of service ensures that the new network infrastructure can be maintained by support staff after the implementation.

Once the current level of service has been analyzed, consider if any additional training is required. This will include training internal IT staff as well as end users.

If technical support has been outsourced, determine if the external company has the skills and knowledge to maintain, support and troubleshoot the new infrastructure.

Analyze Network Requirements

Once the physical network has been analyzed, determine if it is capable of supporting the new infrastructure. This includes LAN/WAN connectivity, available bandwidth, and server distribution.

- Identify physical connections that need to be upgraded to support Active Directory replication and other types of traffic.
- Some types of traffic can be controlled by creating Active Directory sites.
- Identify if any additional servers are required to support the new infrastructure and where on the network they will be located.

Design a Computer and User Authentication Strategy

Identify Common Authentication Requirements

For the system to allow access to resources and data, it must first confirm the identity of the person attempting the access. This identity-confirmation process is called authentication. Basically, there are three different ways of looking at authentication, all based upon what the specific task at hand is.

In order to implement the required authentication methods, determine the authentication requirements of an organization

- Determine the authentication methods used by operating systems and applications. Not all software supports Kerberos.
- If PKI is being used, determine how many certification authorities (CAs) are required or if an external CA will be used.
- Determine if there are enough domain controllers to service all the authentication requests.

Select Authentication Mechanisms

- Local logon and workgroup networking:
 - ▶ NT Lan Manager (NTLM) authentication is used to perform local and workgroup authentication. If a user logs on locally to a workstation or standalone server (in other words, the system is not a member of a domain), NTLM handles the authentication process to ensure that the local username and passwords match. NTLM is pretty straightforward with regard to how it does its job: It looks at the local user account database, finds the requesting user account, and ensures that the passwords match.
- Domain logon and domain networking:
 - ▶ *Kerberos* authentication is used to handle domain logons and network access. Kerberos is used primarily to eliminate the possibility of a network "eavesdropper" from tapping into data over the network, and particularly user names and passwords. The way it works involves an added layer of complexity, and a few new players on the field. Kerberos is discussed in the next section.
- Interforest and Internet authentication:
 - ▶ *Public Key Infrastructure (PKI)* is a set of tools and utilities that are used to ensure secure communications between systems that do not share a common security umbrella, such as membership within the same domain or forest, or Web-based client access. An example of a PKI implementation is a SQL database containing a product catalog that was placed on the Internet and access-restricted to specific personnel. PKI can do this securely without requiring an administrator to create a domain account and password that allows a potential stranger to gain access to other domain resources (such as those to which the Everyone group has access).

See [Access Management](#) for more information.

Kerberos

Kerberos works in the same way as you going to see a movie. First, you go to the ticket counter, tell the person what movie you want to see, and get your ticket. After that, you go to a turnstile or ticket stand and hand the ticket to someone else, and then you're "in." In a nutshell, that's Kerberos.

You won't be able to answer any exam questions based on this description, so the following movie scenario helps you get a good grip on what's going on while learning Kerberos terms:

1. You walk up to a ticket booth. The ticket booth is called the *Key Distribution Center (KDC)*.
2. The KDC has two things in it: a person and a ticket printing machine. The person is called the *Authentication Service*, and the ticket machine is called the *Ticket Granting Service*.
3. You say "I need a ticket." The KDC responds with "What movie?." Actually, the KDC doesn't say this. Instead, his response is encrypted. If you can decrypt it (decryption happens through a password-based algorithm), then you're okay to move on.
4. You tell the Authentication Service, "Star Wars." Again, you're not really saying that. Instead, you're repeating the decrypted response from step 3.
5. The Authentication Service (person) turns to the Ticket Granting Service (ticket printing machine) and gives you a ticket. This ticket is called a *Service Ticket*.
6. You then go to the theater door and present the Service Ticket to the person waiting there. This person is actually the Networking Services on the server that contains the data that you want to see.
7. The person looks at the ticket. If he recognizes it as valid, you gain entry to the theater area. In this case, it's the server.
8. From that point, Kerberos is done. The server itself refers to its internal or replicated user account and password lists to determine appropriate access to the data.

Figure 15 illustrates the process of Kerberos authentication using the movie scenario.

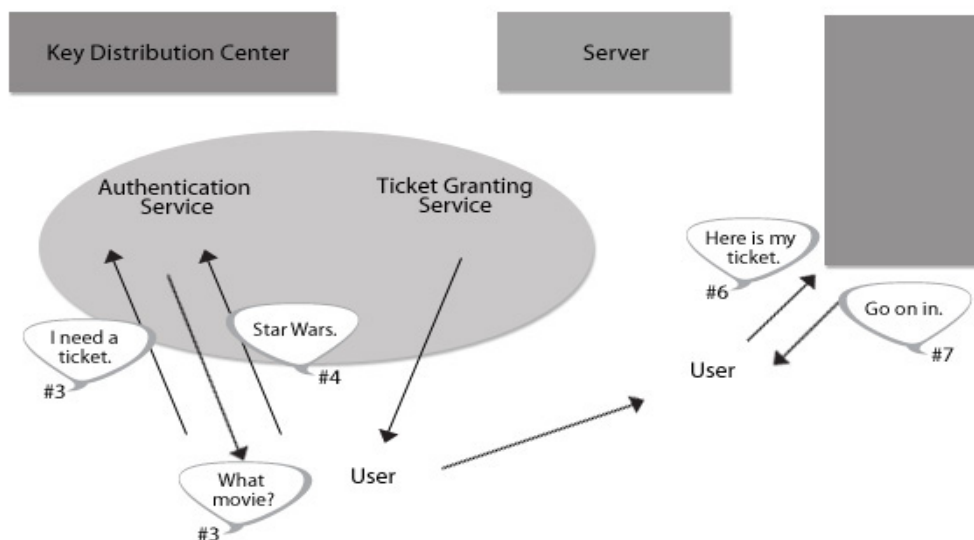


Figure 15 - Kerberos Authentication

As you might imagine, a single KDC can handle tickets for multiple servers, in the same way that a single domain controller can handle authentication for numerous systems in a domain. Kerberos doesn't really need to be configured: it works automatically within any domain.

If you've ever wondered how authentication works between trusted domains, Kerberos is there too. The only difference between an interdomain authentication and standard domain authentication is that the domain KDC maintains a relationship with the trusted/trusting domain's KDC, which basically allows the KDC in one domain to grant tickets for the trusted/trusting domain. Going back to the movie metaphor: it's just like buying a ticket at one movie theater for another theater that belongs to the same corporation.

See [Kerberos V5 authentication](#) for more information.

PKI

- Most secure authentication method
- Requires a PKI infrastructure
- Uses public and private keys to verify the identity of a user or computer
- Implement PKI:
 - ▶ In non-trusted environments
 - ▶ Remote access
 - ▶ Internet access
 - ▶ Computers that do not support Kerberos authentication

Smart Cards

- This authentication method is an extension to PKI and certificates.
- A PIN number is required to access the certificates.
- In multi-factor authentication:
 - ▶ The user must type the correct PIN number to access certificates.
 - ▶ Certificates are required to authenticate to Active Directory.

Design a User and Computer Account Strategy

A user account is an object that is used to represent a user logon to the system. An account contains an account name, and may contain a password, along with other options for additional security.

A computer account is also an object in Active Directory that is used to represent the computers on the network.

You need to identify which individuals within an organization will be granted the rights to create and manage user accounts, to create and manage group accounts, and to create and manage computer accounts.

Specify Account Policy Requirements

- Account policies consist of a password policy, account lockout policy, and Kerberos policy as shown in the figure.

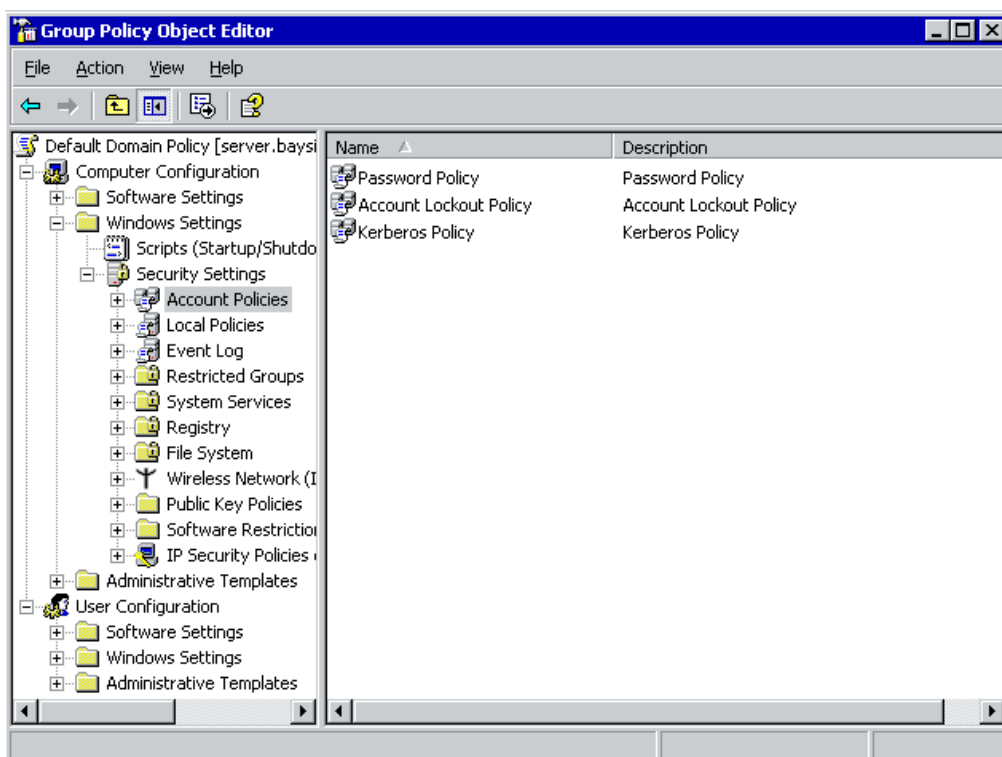


Figure 16 - Account Policies

- Each policy must be applied at the domain level. They can not be applied at the OU level.
- Password policy
 - The password policy determines such things as how often users are required to change their password and the password history.

- Account Lockout policy
 - The account lockout policy defines how the system monitors failed log on attempts and the action to take when a certain number of failed log on requests is reached.
- Kerberos policy
 - The Kerberos policy configures settings such as the maximum lifetime for user and service tickets.

Design Migration Paths to Active Directory

Windows NT 4.0 can be upgraded to Windows Server 2003. Service pack 5 or later is required to upgrade. (See *Table 4*.)

Operating System	Windows Server 2003, Standard Edition	Windows Server 2003, Enterprise Edition	Windows Server 2003, Datacenter Edition
Windows NT Server 4.0	Yes	Yes	No
Windows NT Server 4.0 Enterprise Edition	No	Yes	No
Windows NT Server 4.0 Terminal Server	Yes	Yes	No

Table 4 – Windows NT 4.0 Upgrade Paths

Windows 2000 can be upgraded to Windows Server 2003. (See *Table 5*.)

Operating System	Windows Server 2003, Standard Edition	Windows Server 2003, Enterprise Edition	Windows Server 2003, Datacenter Edition
Windows 2000 Professional	No	No	No
Windows 2000 Server	Yes	Yes	No
Windows 2000 Advanced Server	No	Yes	No
Windows 2000 Datacenter Server	No	No	Yes

Table 5 - Windows 2000 Upgrade Paths

Different migration strategies are available when moving from Windows NT 4.0 or Windows 2000 to Windows Server 2003.

The migration paths include:

- In-place upgrade
- Domain restructuring
- New Active Directory environment

Define Whether Migrations will Include an In-place Upgrade, Domain Restructuring, or Migration to a New Active Directory Environment

- The migration paths are summarized *Table 6*.

Migration Path	Description
In-place upgrade	The existing domain infrastructure is maintained. Change is minimized.
Domain restructure	The existing domain infrastructure is altered. Resource domains may be consolidated into organization units. Resolves any issues with the existing domain model. Migrate user accounts between domains using the Active Directory Migration Tool (ADMT).
New Active Directory environment	A completely new infrastructure is created.

Table 6 - Migration Paths to Windows Server 2003

Design a Strategy for Group Policy Implementation

Group policies are used to control the computing environment. When designing a group policy implementation strategy, you need to consider delegation of administration, the application of GPOs, inheritance modification, and security filtering.

Design the Administration of Group Policy Objects

- Identify who is responsible for administering GPOs.
- Identify each individual's scope of authority over administering GPOs.
- Different ways of organizing GPOs
 - Single policy
 - Separate policies are created for each group of settings such as application settings, desktop settings, and so on
 - Good for decentralized administration
 - Multiple policies
 - One GPO contains all the policy settings that must be applied to a container
 - Good for centralized administration
 - Dedicated policies
 - Policy settings are divided into user settings and computer settings
 - User settings and computer settings are contained in different GPOs
- Administration of GPOs can be divided into different categories:
 - Creating GPOs
 - Modifying GPOs
 - Linking GPOs

Design the Deployment Strategy of GPOs

- Delegation is the process of decentralizing network administration.
 - Assign specific administrative duties to other users or groups within the business.
 - This is another reason for creating organizational units.
- GPOs can be applied locally or at the site, domain, and OU level.
 - GPOs are applied in the following order: local computer, site, domain, OU
 - Where GPOs are linked will determine its scope
 - Link GPOs at the site level to take advantage of physical connections
 - Must be a member of the Enterprise Admins group to link at this level
- GPOs linked at the domain level affect all users and computers in the domain.

- GPOs linked at the OU level provide the finest granularity of control.
 - GPOs can not be linked to containers including the Users and Computers containers.
- Use filtering to exclude certain groups being from affected by a group policy.
 - Remove the Apply Group Policy permission for a user or group to exempt a user or group from the policy settings.
- Enable Block inheritance to prevent inheritance of GPOs.
 - The inheritance of a GPO can be modified so that it is not passed on from parent object to child object.
 - Any policy applied at the site, domain, or OU level can be blocked.
- Enable No Override/Enforced to ensure a GPO can not be blocked.
 - If the option is set, any group policies linked to a parent object will be applied to the child object, regardless of whether the Block Policy Inheritance option is set.

Design a NetBIOS Name Resolution Strategy

In a pure Windows 2000 or Windows Server 2003 environment, NetBIOS may not be required since Active Directory uses DNS. NetBIOS is supported by Windows Server 2003 to remain backwards compatible.

- Identify if there are any computers or applications that rely on NetBIOS.
- NetBIOS names can be resolved through broadcasts (not recommended), LMHOSTS files, and WINS.
- LMHOSTS files:
 - Are static text file that maps NetBIOS names to IP addresses.
 - Must be manually updated by an administrator.
 - Use the #pre directive to preload specific entries into the NetBIOS name cache.
- WINS is still used in Windows Server 2003.
 - WINS provides NetBIOS name resolution for legacy clients and applications.
 - WINS allows name resolution across routers.
 - WINS will be needed until all legacy clients and applications are removed from the network.
- When upgrading from Windows NT 4.0, determine if the existing NetBIOS name is suitable for the DNS prefix.
- It's a good practice to make the NetBIOS name and the DNS name be identical. This results in less confusion for technical staff and end users as well.

Design a WINS Replication Strategy

- WINS servers can replicate their databases.
 - Improves name resolution
 - Fault tolerance through secondary WINS servers
- Replication can occur at a specific interval and when a certain number of database changes have occurred, as shown in *Figure 17*.

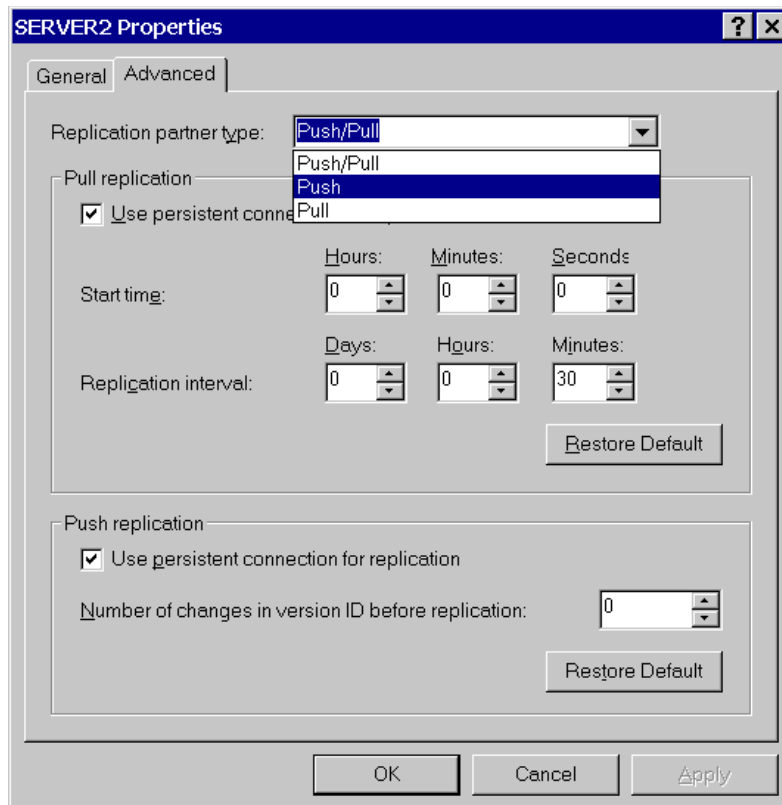


Figure 17 - WINS Replication

- The three replication options include:
 - ▶ Push
 - ▶ Pull
 - ▶ Push/Pull
- Push replication occurs between two WINS servers after a specified number of changes on a database.
 - ▶ Ideal option if you want to minimize traffic
 - ▶ Should not be used alone if you want the WINS databases to be very accurate
- Pull replication occurs on a specific interval, regardless of the number of changes.
 - ▶ Configure WINS servers to replicate when the network is not busy
- Push/pull is a combination of the push and pull options. Replication occurs when a certain number of changes occur and at a specific time.
- Convergence is the amount of time for changes to be replicated to other WINS servers on the network.
- Calculate the convergence time by adding the replication interval times.

Design a Remote Access Strategy

Specify the Remote Access Method

- Remote access supports [dial-up connections](#), [VPN connections](#), and wireless connections.
- Dial-up connections:
 - ▶ Select this connection method when connecting over the Internet is not acceptable.
 - ▶ Added security is achieved through callback and caller identification.
 - ▶ You can configure user dial-in settings through a user account properties dialog box.
 - From the Dial-in tab, you can configure several other settings including caller ID, callback options, and static IP routes.
 - If you configure the settings for the user account, they must match the settings configured on the client or the connection attempt will be denied.
- Windows Server 2003 supports PPP and SLIP for dial-up connections.
 - ▶ Windows Server 2003 only supports SLIP for dialing out.
- VPN connections:
 - ▶ Use a public network, such as the Internet, as a backbone.
 - ▶ Reduce remote access connection costs.
- The L2TP and PPTP tunneling protocols are used for VPNs.
- PPTP is older than L2TP.
 - ▶ Windows 2000 Professional and later clients support L2TP natively.
 - ▶ Windows 98 and Windows NT have an L2TP client that can be downloaded from Microsoft.
- Wireless clients connect to a network by using radio frequencies ranging from 2.4 GHz to 5.0 GHz.

Specify the Authentication Method for Remote Access

- Authentication protocols include:
 - ▶ [Password Authentication Protocol \(PAP\)](#) - The least secure of all authentication protocols because it sends the user name and password in clear text.
 - ▶ [Shiva PAP \(SPAP\)](#) - Can be used to authenticate against Shiva remote access servers and to authenticate against Windows 2003 Servers. This protocol is typically more secure than PAP but not as secure as CHAP or MS-CHAP.
 - ▶ [Challenge Handshake Authentication Protocol \(CHAP\)](#) - Does not send the user name and password across the network. Instead, it uses a challenge response with a one-way hash algorithm. It is an industry-standard protocol that can be used to authenticate non Windows-based clients.
 - ▶ [MS-CHAP versions 1 and 2](#) - A Microsoft version of CHAP, MS-CHAP uses mutual authentication and encryption for Windows-based clients. MS-CHAP version 2 provides strong encryption and separate encryption keys for sending and receiving data.
 - ▶ [Extensible Authentication Protocol \(EAP\)](#) - An extension of the PPP protocol that provides support for other authentication mechanisms, such as smart cards. This authentication protocol requires the presence of a PK infrastructure.

- MS-CHAP v2 and EAP are recommended.
- Two types of encryption are MPPE and IPSec.
 - MPPE uses 40-bit, 56-bit, and 128-bit. Some legacy clients do not support 56-bit.

See [Enable authentication protocols](#) and [Configure identity authentication and data encryption settings](#) for more information.

Design the Remote Access Infrastructure

Plan Capacity

- It is important that users have access their folders and files that are stored on the network.
- Users can gain access to network resources from remote locations using remote access.
- Remote access servers are the link between remote access users and the internal network.
- With the proper credentials and connection settings, a user can gain remote access to network resources.
- Multiple remote access servers may be required.
 - Use IAS to store remote access policies in a central location.
 - IAS centralizes accounting, authentication, authorizations, and auditing of remote access connections.
- A RADIUS proxy forwards remote access connection requests based on configured rules.
- Consider placing a second remote access server on each subnet that already contains a remote access server.
- Distribute remote access servers throughout different business locations.

Design Security for Remote Access Users

Identify Security Host Requirements

Many businesses permit their users and clients to work outside of the office. Using remote access users can gain access to network resources from any location. One of the challenges IT users face is providing remote access while still maintaining a high level of security.

Identify the Authentication and Accounting Provider

- Internet Authentication Service (IAS) centralizes authentication and logging for multiple RAS servers. In regard to IAS services, the RAS servers are the client of an IAS server.

Design Remote Access Policies

- [Remote access policies](#) define when, where, and how a user can gain access to a network from a remote location.
- They contain the elements of conditions, permissions, and profile.
- Permission options for user accounts include: **Allow access**, **Deny access**, and **Control access through Remote Access Policy**.
- The **Control access through Remote Access Policy** option is unavailable on the Dial-in tab if the domain functional level is set to Mixed.
- The default remote access policy (**Connections to routing and remote access server**) denies remote access.
- The permissions in a remote access policy include: **Grant remote access permission** and **Deny remote access permission**.
- A connection attempt must match the settings within a remote access policy.
- Policies are stored locally on the remote access server, not in Active Directory.

Conditions

- Remote access [conditions](#) are used to place restrictions on the initial connection attempt by a remote access client.
- When a user attempts to access a remote access server, the properties of the connection attempt must match the conditions of at least one remote access policy.

Permissions

- Remote access [permissions](#) are similar in function to permissions you set on a folder or file.
- They determine which users are permitted or denied remote access. Remote access permissions can be configured through the properties of a user account or through the remote access policy.
- Remote access permissions configured from the properties of a user account override those configured in a remote access policy.

Profile

- The [profile](#) of a policy determines the properties that a connection attempt must match.
- Once a connection attempt has been authorized, meaning the remote access user has the required permission, the profile settings are then applied to the connection attempt.

Remote Access Policy Evaluation

1. A user attempts to establish a connection. The connection attempt is evaluated against the conditions of the remote access policy. If multiple policies exist, the first policy in the list is evaluated.
2. If the conditions do not meet the conditions of the first policy, the conditions of each policy are evaluated until a match is found.
3. If the connection attempt does not match the conditions of any policy, it is rejected. If the connection attempt does match the conditions of a policy, the evaluation process continues and the permissions of this policy are evaluate.
4. If the user's account properties grant access, the profile settings are evaluated. If the user's account properties deny access, the connection attempt is rejected.
5. If access is controlled through the policy permission, access is either granted or denied based on the permission settings.
6. If the user's account properties grant access or if the policy permissions grant access, the profile settings of the policy and the properties of the user account are evaluated.
7. If the connection attempt matches both the account settings and the policy profile settings, the connection attempt is granted.
8. If the connection attempt does not match the account settings and the policy profile settings, the connection attempt is rejected.

Specify Logging and Auditing Settings

- Use IAS if there are multiple remote access servers.
- A server running IAS is called a [RADIUS server](#).
- Remote access polices are stored in a central location (on the IAS server).
- IAS servers also log remote access connections.
- See [Logging](#) for more information.

Design a DNS Service Implementation

Design a Strategy for DNS Zone Storage

- To host an Active Directory-integrated zone, a DNS server must also be a domain controller.
- Incremental Zone Transfer (IXFR) allows for more frequent [zone transfer](#) and thereby increases accuracy of zone information.
- Caching-only DNS servers should be used with a small remote office that has a relatively slow link back to the main office. These servers are not authoritative for a zone and therefore do not perform zone transfer. This conserves available bandwidth on the slow link.
- Conditional forwarding is new to Windows Server 2003. This service allows all queries for a particular namespace to be forwarded directly to the server that hosts that namespace. Doing so increases the efficiency of name resolution on a network.
- A zone is a discrete, contiguous portion of a DNS namespace. Zone information is contained in a zone database file. Adding the words database file to the word zone sometimes helps tremendously when answering a question relating to zones.
- A standard primary zone database file is the original database of a zone. It can be written to and read from as well.
- A standard secondary zone database file is a copy of the original database for a zone. It can be read but cannot be written to, except through zone transfer. Standard secondary zones are used to help offload some of the traffic between DNS servers.
- Servers that host a standard primary zone or secondary zone can be master servers for other servers that host the standard secondary database for that zone.
- Zone transfer can be bandwidth-intensive, especially on slow links.
- An important part of DNS design is understanding and working within bandwidth limitations in regard to zones.
- A Standard Secondary zone database file does provide some short-term fault tolerance for name resolution but cannot be considered a long-term fault-tolerance option.
 - For a long-term fault-tolerance solution, implement Active Directory-integrated zones.
- Active Directory-integrated zones are preferred when minimizing Zone Transfer is a key concern. Because the replication between Active Directory-integrated zones occurs with Active Directory replication, there is said to be no zone transfer at all with Active Directory-integrated zones.
- DHCP can interoperate with DNS and thereby register hostnames of clients in the DNS database. Clients that can register their own A (host) records include Windows 2000 Professional and Windows XP Professional.
- Active Directory-integrated zones are preferred when available.
 - Active Directory-integrated zones do not use zone transfer, but instead piggy-back their zone change information on Active Directory replication.
 - Active Directory zones are all primary.
 - Active Directory zones can be secure or not secure.
- Stub zones are special zones that contain only the SOA resource record, name server record, and glue A host record (IP address) for the zone. These zones are used in networks with noncontiguous namespaces to make DNS name resolution more efficient.

Specify the Use of DNS Server Options

- Several different [DNS server options](#) can be configured from a DNS server's properties dialog box.
 - ▶ *General tab* - Configure the interfaces on which the DNS server will listen for DNS queries.
 - ▶ *Forwarders tab* - Configure where a DNS server can forward DNS queries to that it cannot resolve.
 - ▶ *Advanced tab* - Configure advanced options, determine the method of name checking, and determine the location from which zone data is loaded, as well as enable automatic scavenging of stale records.
 - ▶ *Root Hints tab* - Configure root name servers that the DNS server can use and refer to when resolving queries.
 - ▶ *Debugging tab* - Enable debugging. When this option is enabled, packets sent and received by the DNS server are recorded in a log file. You can also configure the type of information to record in the file.
 - ▶ *Event Logging tab* - Configure the type of events that should be written to the DNS event log. You can log errors, errors and warnings, and all events. You can also turn off logging by selecting No Events.
 - ▶ *Monitoring tab* - Use to verify the configuration by manually sending queries against the server. You can perform a simple query that uses the DNS client on the local server to query the DNS service to return the best possible answer. You can also perform a recursive query in which the local DNS server can query other DNS servers resolve the query.

Identify the Registration Requirements of Specific DNS Records

- [Resource records](#) are created within the DNS management console.
 - ▶ A records map host names to IP addresses.
 - ▶ MX records route messages to a specific mail exchanger.
 - ▶ PTR records map IP addresses to host names.
 - ▶ CNAME records create aliases for names that are already referenced by another resource record.
- Stub zones contain only certain resource records:
 - ▶ SOA
 - ▶ NS
 - ▶ A

Specify the Server Specifications to Meet System Requirements

- There are four editions of Windows Server 2003: Web, Standard, Enterprise, and Datacenter. Web Edition cannot be used as a domain controller.
- The Windows Server 2003 platforms are versatile in their design and can therefore perform in a variety of network environments and function in a variety of different server roles including:
 - ▶ Mail server
 - ▶ File and print server
 - ▶ Web server
 - ▶ Application server
 - ▶ RAS/VPN server
 - ▶ Terminal server
 - ▶ DHCP/WINS/DNS Server
 - ▶ Domain Controller
- They are all designed to provide dependability, productivity, security, and connectivity through their standard features.
- Each platform has unique features and has been built with different network environments in mind.

Web Server

- Windows Server 2003 Web Server is specifically designed for those small organizations or departments needing to build and host web applications and web services.
- The functionality behind Web Server is targeted towards ISPs and developers to provide web services.
- The technology of Web Server incorporates ASP and the .NET framework allowing web developers to easily build and deploy web applications and services.

Web Servers functionality is not intended to go beyond that of a stand alone web server and does not provide the full functionality of a server operating system. For example, while Web Server can be a member of a domain, it cannot run Active Directory. So for organizations looking for more functionality from their server operating system, one of the other Windows Server 2003 platforms should be considered.

Standard Server

- Windows Server 2003 Standard Server is designed to meet the needs of smaller organizations.
- It provides a high level of availability and all the features needed to meet the needs of work-groups, branch offices, and departments.
- Its offers solutions for file and printer sharing, centralized management through Active Directory, secure Internet connectivity, and Web services.
- It offers a certain level of scalability supporting two-way symmetric multi-processing and up to 4 GB of RAM.

For organizations needing to achieve the highest level of availability, Enterprise Server or Datacenter Server is the ideal choice.

Enterprise Server

Many organizations today are running critical line-of-business applications, messaging applications, and e-commerce websites that are crucial to their day-to-day operations. In other words, these applications and services need to be up and running. Any downtime or decrease in performance can in the end mean a loss of business or revenue. The features and capabilities of Windows Server 2003 Enterprise Server help organizations address the issue of high-availability and scalability.

- Windows Server 2003 Enterprise Server is designed for medium to large organizations.
- It has all the features of Standard Server with additional features targeted for businesses requiring a high level of availability and the ability to scale applications up and out as demands increase.
- Scalability or “scaling up” can be achieved by adding more horse power to a server as demands increase.
 - ▶ One of the ways in which Enterprise Server addresses scalability is through Enhanced Symmetric Multiprocessing (SMP).
 - ▶ To increase server performance, you can add up to 8 processors.
 - ▶ Enterprise Server includes enhanced memory capabilities, allowing you to add up to 8 GB of memory.
 - ▶ To provide high availability and fault tolerance for mission-critical applications, services, and data, Enterprise Server supports server clusters through Cluster Service.
 - ▶ Cluster Service provides a means of connecting multiple servers, up to 8 nodes, that operate and appear as a single system.
 - ▶ Cluster Service provides high availability by automatically detecting when an application or service fails and automatically restarting it on another cluster member.
 - ▶ Downtime from failed applications and services can be as little as a few seconds and go unnoticed by users.
- Enterprise Server and Datacenter Server both support clustering. This feature is not available in Standard Server or Web Server.

Datacenter Server

- Windows Server 2003, Datacenter Server is designed for large enterprise organizations requiring the highest level of server performance.
- It's by far the most powerful operating system out of the four Windows Server 2003 platforms, offering the highest level of scalability, reliability, and availability.
- Unlike the other version of Windows Server 2003, Datacenter Server cannot be purchased on its own.
- Some of the features that set Datacenter Server apart from the other operating system versions are its support for 32-way symmetric multi-processing and up to 64 GB of RAM (whereas the 64-bit version is capable of supporting 64 processors and 128 GB of RAM).

Hardware Requirements

Table 7 outlines the minimum and recommended hardware requirements for each of the platforms.

Platform	Hardware Requirements
Standard Server	133 MHz processor (550 MHz recommended) 128 MB of RAM (256 MB recommended) 1.5 GB of free disk space
Web Server	133 MHz processor (550 MHz recommended) 128 MB of RAM (256 MB recommended) 1.5 GB free disk space
Enterprise Server	133 MHz processor (733 MHz recommended) 128 MB of RAM (256 MB recommended) 1.5 GB free disk space
Datacenter Server	400 MHz processor (733 MHz recommended) 512 MB of RAM (1GB recommended) 1.5 GB free disk space

Table 7 - Minimum and Recommended Requirements for Windows Server 2003 Platforms

- When planning for server hardware, always consider applications, services, and the server role.

Design Internet Connectivity

- [Network Address Translation \(NAT\)](#):
 - ▶ Provides a connectivity solution for non-routed networks
 - ▶ Includes a DHCP allocator
 - ▶ Provides a name resolution component
 - ▶ Uses static mappings to make private resources accessible to Internet users
 - ▶ Provides a basic firewall to protect internal network from unsolicited traffic
 - ▶ NAT-T provides support for L2TP/IPSec VPN connections across a NAT interface
 - ▶ Use with VPN connections for user-level authentication
- [Internet Connection Sharing \(ICS\)](#):
 - ▶ Provides a connectivity solution for small networks
 - ▶ Ideal for networks with less than 10 workstations
- Proxy server:
 - ▶ More secure than NAT and ICS
 - ▶ Internet connectivity based on protocols, users, groups, and applications

- Firewall:
 - Provides packet filtering
 - Serves as a bastion host

Design a Network and Routing Topology for a Company

Design a TCP/IP Addressing Scheme Through the Use of IP Subnets

- IP addresses:
 - Static:
 - Servers
 - Network printers
 - Router interfaces
 - Automatic Private IP addressing
 - IP address in the range of 169.254.0.1 - 169.254.255.255
 - If a DHCP server is unavailable, Automatic Private IP Addressing is used.
 - If using APIPA, a computer will not be able to communicate outside of the local subnet.
 - DHCP
- IP address scheme
 - Public IP address
 - Suitable for small networks
 - Each computer is assigned a public IP address
 - Less secure than private IP addressing
 - Private IP address
 - 10.0.0.0/8, 172.16.0.0/16, 192.168.0.0/24
 - Internet access can be provided through a proxy or NAT server
 - More secure and flexible than public IP addressing
- Public IP addresses should be registered through the Internet Corporation of Assigned Names and Numbers (ICANN) or one of its registries.

Specify the Placement of Routers

- Routers use the information within an IP packet header to determine the destination IP address.
- Routers maintain information within a table about the physical network, such as the path to a destination network and the metric associated with the route, the metric being the distance between the source and destination networks.

Design IP Address Assignment by Using DHCP

- DHCP server leases IP addresses to DHCP clients.
- Can assign other parameters to clients, as shown in *Figure 18*.
- Server options can be configured at the server, scope, class, or client level:
 - ▶ Options are applied in the following order: server, scope, class, client.
 - ▶ Options configured at the server level are applied to all DHCP clients.
 - ▶ Options configured at the scope level are applied only to DHCP clients on a specific subnet.
 - ▶ Options configured at the class level are applied to a group of users or computers with similar needs or vendor information.
 - ▶ Options configured at the client level are applied only to the specific DHCP client.
- Types of server options:
 - ▶ 006 DNS server
 - ▶ 003 Router
 - ▶ 015 DNS Domain Name
 - ▶ 044 WINS\NBNS Servers
 - ▶ 046 WINS\NBT Node Type
- It is a best practice to implement at least two DHCP servers for fault tolerance.
 - ▶ Increase the availability of DHCP using the 80/20 rule.
- The number of DHCP servers required will depend on the network configuration, number of DHCP clients, routing requirements, and the hardware being used.
- If using multiple DHCP servers, consider placing one on either side of a slow WAN link or dial-up connection.
 - ▶ DHCP requests do not need to be sent across slow connections.
- Most routers block DHCP broadcasts
 - ▶ A DHCP relay agent is required to forward DHCP broadcasts between subnets.
 - ▶ You can configure a router or a workstation to be a DHCP relay agent.
- Integrate DHCP with Active Directory to prevent rogue DHCP servers on the network
 - ▶ All DHCP servers must be authorized in Active Directory before they can lease IP addresses.

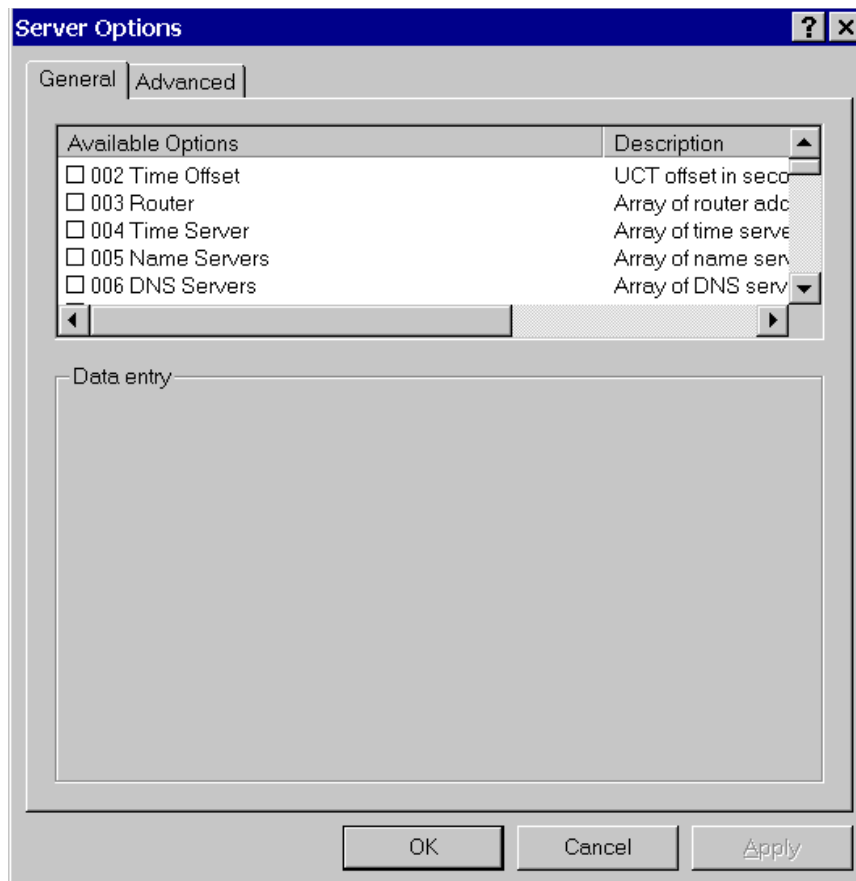


Figure 18 - Optional Parameters

Design a Perimeter Network

- Many organizations use a combination of firewalls to create an area of the network that is neither completely on the inside, nor completely on the outside.
- Physically placing a server between the Internet and the other servers on a private network can provide additional authentication and security for a network while letting the other servers keep more services available for those who are authorized.
- A perimeter network is a small network that sits between the Internet and the private network.
- Firewalls are used to protect the private network from intrusion.
- Servers that are sometimes placed into perimeter networks include DNS, Web, FTP, e-mail, and others.
- Typically, servers that have vital information about the internal network, such as Domain Controllers, are not placed into a perimeter networks.

Practice Questions

Case Study 1

Background

MVP Inc. is a large clothing company with 150 retail stores throughout the U.S. and Canada. Ten retail stores are located in Europe and Japan. The headquarters office is located in Portland. There are four distribution centers in Chicago, Los Angeles, Tampa, and Toronto.

Current LAN/WAN Infrastructure

Four resource domains are in place for each of the distribution centers. The master account domain is located in Portland. Domain controllers run Windows NT 4.0. Primary domain controllers (PDCs) are located in the headquarters office and each distribution center. Each distribution center also hosts a backup domain controller (BDC) for the master account domain.

DNS servers run BIND version 8.1.2. The organization currently has no plans to upgrade them.

Workstations run a variety of platforms including Windows 95, Windows 98, and Windows 2000 Professional.

The headquarters office uses a 100-Mbps Ethernet network. Each distribution center uses a 10-Mbps Ethernet network. There are no plans to upgrade the existing LAN structure.

Each distribution center is connected to the Portland office using a 56-Kbps frame relay connection. Retail locations are connected to the closest distribution center using dial-up connections.

Proposed LAN/WAN Infrastructure

The company plans to migrate to Windows Server 2003. There are no plans to upgrade the existing LAN or WAN structure. The following issues should be addressed by the new infrastructure:

- Decrease name resolution response time for users within distribution centers without increasing network traffic on the WAN connections.
- Obtain IP addresses from a DHCP server.
- Register all client computers dynamically with a DNS server.
- Provide name resolution services for all clients.
- Replication should occur during non-business hours.

Current Internet Positioning

MVP Inc. currently has an Internet presence. The company has registered the domain name mvp.com. A subsidiary company has recently been created called MVP Kids.

Active Directory/Network Services Commentary

IT Manager: Staff members in the IT group spend too much time dealing with minor issues such as creating user accounts and changing passwords, among others. I would like the new infrastructure to allow us to delegate some of these tasks to other administrators.

Retail Store Managers: The process for creating a user account is slow. We fax new account information to the headquarters office and then wait for several days before the user account is created. The process should be more efficient.

CIO: I am not satisfied with the existing infrastructure. I would like some of the existing domains to be consolidated.

1. Will the current DNS infrastructure support Active Directory?
Select the best answer.
 - A. No.Active Directory only supports Windows Server 2003 DNS.
 - B. Yes.BIND 8.1.2 fully integrates with Windows Server 2003.
 - C. Yes.BIND 8.1.2 supports SRV records.Resources records will have to be added manually because this version of BIND does not support dynamic updates.
 - D. No.BIND 8.1.2 does not support SRV records or dynamic updates.

2. How many sites should be created for MVP Inc.?
Select the best answer.
 - A. One site is sufficient for the entire company.
 - B. Two sites --one for the headquarters office and one for the distribution centers.
 - C. Four sites --one for each of the distribution centers.
 - D. Five sites --one for the headquarters office and one for each of the distribution offices.

3. John is designing the connectivity between sites.How should he configure the site links between the company's headquarters office and the distribution centers?
Select the best answer.
 - A. Create site links that use SMTP and configure a replication schedule so replication occurs only during non-business hours.
 - B. Create site links that use SMTP.Do not configure a replication schedule.
 - C. Create site links that use RPC over IP and configure a replication schedule so replication occurs only during non-business hours.
 - D. Create site links that use RPC over IP.Do not configure a replication schedule.

4. An administrator in the Toronto distribution office makes a change to the properties of a user account.How long will it take for this change to be replicated to other domain controllers within Toronto?
Select the best answer.
 - A. 5 seconds
 - B. 30 seconds
 - C. 5 minutes
 - D. 30 minutes
 - E. 60 minutes

5. Which of the following will impact the number of domains created for the Active Directory infrastructure?
Select the best answer.
- A. The requirement of unique security policies throughout the company
 - B. The speed of the WAN connections between the headquarters office and distributions centers
 - C. The company's existing Internet presence
 - D. The administrative model currently implemented by the company
6. How many domains should be created for MVP Inc.?
Select the best answer.
- A. One domain for the entire organization
 - B. Two domains --one for the headquarters office and one for the distribution offices
 - C. Four domains --one domain for each of the distribution centers
 - D. Five domains --one for the headquarters office and one for each of the distribution centers
7. Users in the subsidiary, MVP Kids, require a specialized application that will make modifications to the schema. This change should not affect users in the MVP Inc. organization. What should you do?
Select the best answer.
- A. Configure a separate OU within the existing forest for MVP Kids.
 - B. Configure a child domain within the existing forest for MVP Kids.
 - C. Configure MVP Kids in its own tree within the existing forest.
 - D. Configure a new forest for MVP Kids.
8. Users in MVP Inc. need access to resources within MVP Kids and vice versa. How can this goal be accomplished?
Select the best answer.
- A. Two-way transitive trusts are automatically created between forest root domains. Therefore, no further configuration is required.
 - B. No additional configuration is required. A shortcut trust is automatically configured between forest root domains.
 - C. Configure a one-way NTLM trust between the forest root domains.
 - D. Configure a forest trust between the two forest root domains.

9. Jim is a member of the Active Directory and Network Services design team. He is designing the Active Directory and network infrastructure. He wants to increase the availability of various services within each distribution center. How should he proceed?
Select four.
- A. Place a schema master within each of the distribution offices.
 - B. Place a Global Catalog Server within each distribution office.
 - C. Place a domain naming master in each distribution office.
 - D. Place a domain controller within each distribution office.
 - E. Place a DNS server within each location.
 - F. Place a WINS server within each location.
10. You are analyzing the existing domain model to prepare for the migration to Windows Server 2003. What is the minimum number of trusts that exist in the current environment before the upgrade?
Select the best answer.
- A. One
 - B. Two
 - C. Four
 - D. Eight
11. MVP Inc. is implementing DHCP for assigning IP addresses. The DHCP server will also assign various other options to DHCP clients. Each distribution center hosts a router, so the 003 router option will be configured on the DHCP server. When configuring this option, what should you do?
Select the best answer.
- A. Configure the 003 router option at the server level.
 - B. Configure the 003 router option at the scope level.
 - C. Configure the 003 router option at the class level.
 - D. Configure the 003 router option at the client level.
12. MVP is opening a remote distribution office for MVP Kids. The distribution office will contain a maximum of 30 users. The IT Manager wants to implement a DNS solution that requires minimal administration for the new office without generating additional zone transfer traffic on the WAN connection. How can this goal be accomplished?
Select the best answer.
- A. Configure the HOSTS files.
 - B. Configure a secondary DNS server in the new office.
 - C. Configure a caching-only DNS server in the new office.
 - D. Configure a WNS server in the new office.

13. A network administrator within one of the distribution offices makes a change to the properties of an existing user account. Several days later he notices that the change has not been updated within any of the groups the user is a member of. What is causing the problem?
Select the best answer.
- A. The schema master is not available.
 - B. The RID master is not available.
 - C. The infrastructure master is not available.
 - D. The PDC emulator is not available.
14. When planning the migration from the existing Windows NT 4.0 infrastructure to Windows Server 2003, which of the following statements is correct?
Select the best answer.
- A. The resource domains should be migrated to Windows Server 2003 domains.
 - B. The resource domains should be replaced with OUs for delegation of authority.
 - C. OUs should replace the resource domains to maintain the existing security boundaries.
 - D. Replace the existing resource domain with Windows Server 2003 domains to implement delegation of authority.
15. Jim is migrating the existing infrastructure to Windows Server 2003. During the migration process, the existing PDCs are upgraded to Windows Server 2003. The network temporarily hosts Windows Server 2003 domain controllers and Windows NT 4.0 BDCs. The domain controller functioning as the PDC emulator fails. What should Jim do?
Select the best answer.
- A. Use the Active Directory User and Computers console to seize the role.
 - B. Use the gpupdate command to seize the role.
 - C. Use the Active Directory Users and Computers console to transfer the role.
 - D. Use Ntdsutil to seize the role.

Case Study 2

Background

Contoso Ltd. is a large manufacturing company. The company consists of two distinct divisions. One division, FreeSpace, manufactures a line of home furniture. In 2002, Contoso Ltd. acquired another company, Durance, which manufactures a line of home appliances. Although the companies have consolidated, there still remains a distinction between the two divisions. Each division has its own CIO, who reports to the CEO of Contoso Ltd.

Contoso Ltd. has manufacturing plants in Vancouver, Edmonton, Calgary, and Toronto. The corporate headquarters office is located in Winnipeg. All plants currently manufacture furniture and appliances.

Current LAN/WAN Environment

All LANs currently operate at 100 Mbps. The current WAN structure is shown in the exhibit. Contoso Ltd. is upgrading from Windows NT 4.0 to Windows Server 2003. The company currently implements two domains, one for each division.

Proposed LAN/WAN Environment

There are no proposed changes to the existing WAN environment. There are plans to open a new division under the FreeSpace name for office furnishings. The new infrastructure should meet the following criteria:


- Separate external and internal namespaces
- Reduced network traffic on WAN connections, wherever possible
- Secure zone transfers
- Dynamic IP addressing
- Optimized authentication
- Quick searches for network resources by users

Active Directory/Network Commentary

CEO: The CIOs for each division will maintain complete decision-making authority for their own division. Each will have an annual budget. To maintain a clear distinction between the two divisions, I want the new Active Directory infrastructure to consist of two forests.

CIO: Contoso Ltd. currently has a Web site that clients and customers can access to find information about various products for both divisions. There is a single Web site for both home furnishings and appliances. The company has registered the name of contoso.com.

IT Managers: We would like the ability to delegate some of the administrative tasks to local administrators. However, these local administrators should have permission to perform only specific administrative tasks within their own location.

Exhibit(s):

1. Which of the following factors will have an impact on the forest design for the company?
Select the best answer.
 - A. The speed of the existing WAN connections
 - B. The total number of users within the company
 - C. The administrative model of the company
 - D. The existing DNS implementation

2. Which of the following are technical ramifications of creating two forests?
Select two.
 - A. There is no authentication between forests.
 - B. There will be a single Global Catalog for the entire company.
 - C. Resources cannot be shared between the two forests.
 - D. Schema updates must be performed twice.

3. The company has just finished upgrading to Windows Server 2003. Resources must be shared between all domains in each forest. Universal groups are created in each forest root domain. Bob, a member of the Enterprise Admins group, tries to create a forest trust but is unsuccessful. What is most likely causing this problem to occur?
Select the best answer.
 - A. The forest functional level has not been raised.
 - B. Resources cannot be shared between forests.
 - C. The functional level of the forest root domains is Windows 2000.
 - D. Bob is not a member of the Schema Admins group.

4. Mike is a member of the design team hired by Contoso to plan the upgrade to Windows Server 2003. He has completed the site design for the new infrastructure. He is planning the placement of domain controllers. What is the minimum number of domain controllers that should be deployed? Select the best answer.
- A. 1
 - B. 2
 - C. 5
 - D. 10
5. You need to design a domain name structure for the company. Which of the following should you use? Select two.
- A. Contoso.com
 - B. Freespace.com
 - C. Ad.contoso.com
 - D. Durance.com
6. You are creating the trust design for Contoso Ltd. Which of the following trust relationships are required? Select the best answer.
- A. A one-way trust in which freespace.com trusts durance.com
 - B. A forest trust between freespace.com and durance.com
 - C. A one-way trust where durance.com trusts freespace.com
 - D. A shortcut trust between freespace.com and durance.com
7. Which of the following factors will influence the site design for Contoso Ltd.? Select two.
- A. The number of locations
 - B. The platforms installed on client computers
 - C. The current domain model used by the company
 - D. The number of domain controllers
 - E. The available bandwidth for each WAN connection
8. You need to minimize the impact of replication traffic on the WAN connections. What should you do? Select the best answer.
- A. Create Active Directory integrated zones.
 - B. Create GPOs at the OU level instead of the domain level.
 - C. Create global groups instead of domain local groups.
 - D. Create IP site links between sites.

9. Which of the following site designs should you use for Contoso Ltd?
Select the best answer.
- A. Create one site for all locations.
 - B. Upgrade all WAN connections to 1.5 Mbps. Create a single site for all locations.
 - C. Create one site for each location. Use IP site links to schedule replication.
 - D. Upgrade all WAN connections to 1.5 Mbps. Create a single site for each locations. Use IP site links to schedule replication.
10. The site link cost between Vancouver and Edmonton is set to the default value. What cost value should be set for the link between Vancouver and Calgary?
Select the best answer.
- A. 1
 - B. 10
 - C. 50
 - D. 150
11. The Toronto location requires an additional domain controller. You want to configure an existing member server that is already running Windows Server 2003 as a domain controller. What should you do?
Select the best answer.
- A. Back up all data. Format the hard disk and install Windows Server 2003. During the installation select the Domain Controller option.
 - B. Back up all data. Use the DCPROMO command to upgrade the computer to a domain controller.
 - C. Back up all data. Reinstall the operating system.
 - D. Insert the Windows Server 2003 installation CD into the CD-ROM drive. Select the option to upgrade to a domain controller.
12. Assuming you used the default installation, how many sites will initially exist for Contoso Ltd.?
Select the best answer.
- A. 1
 - B. 2
 - C. 4
 - D. 5

13. Sites and site links are created for each location to connect them to the corporate headquarters office. All of the default values are accepted when the site links are created. Which of the following should you change for the site link connecting Vancouver and Calgary?
Select the best answer.
- A. Site link name
 - B. Cost
 - C. Replication protocol
 - D. Interval
14. Many of the existing servers within Contoso Ltd. will be upgraded to Windows Server 2003. Which of the following platforms can be directly upgraded to Windows Server 2003 Standard Edition?
Select two.
- A. Windows NT Server 4.0
 - B. Windows 2000 Professional
 - C. Windows 2000 Server
 - D. Windows NT Server 4.0 Enterprise Edition
15. Contoso Ltd. is opening a small sales office. The office will have approximately 10 computers running Windows XP Professional. The computers will connect to the headquarters office via a virtual private network. All computers must share a single DSL connection. The Internet connectivity solution must be cost effective and relatively simple to administer. Which technology should you use?
Select two.
- A. NAT
 - B. ICS
 - C. Proxy server
 - D. IAS

Answers and Explanations

Case Study 1

1. Answer: B

Explanation A. BIND 8.1.2 integrates with Windows Server 2003 Active Directory. It supports both SRV records and dynamic updates.

Explanation B. BIND 8.1.2 integrates with Windows Server 2003 Active Directory. It supports both SRV records and dynamic updates.

Explanation C. BIND 8.1.2 integrates with Windows Server 2003 Active Directory. It supports both SRV records and dynamic updates.

Explanation D. BIND 8.1.2 integrates with Windows Server 2003 Active Directory. It supports both SRV records and dynamic updates.

2. Answer: D

Explanation A. The 56-Kbps frame relay links do not provide the necessary connections to create a single site.

Explanation B. The 56-Kbps frame relay links do not provide the necessary connections to place the distribution centers into a single site.

Explanation C. The headquarters office must also be included in the site design. Therefore, five sites would be required.

Explanation D. Sites are created to control replication. Because a 56-Kbps frame relay connection is considered slow, five sites should be created: one for the headquarters office and one for each distribution center. The sites are required because the offices are connected by WAN links.

3. Answer: B

Explanation A. SMTP ignores any replication schedules configured for a site link.

Explanation B. SMTP ignores any replication schedules configured for a site link. SMTP uses a store and forward method of replicating information.

Explanation C. RPC would not be the best replication protocol as it is unreliable over connections that are this slow. If RPC times out, replication will fail.

Explanation D. RPC would not be the best replication protocol as it is unreliable over connections that are this slow.

4. Answer: C

Explanation A. Intra-site replication is configured to occur every 5 minutes by default.

Explanation B. Intra-site replication is configured to occur every 5 minutes by default.

Explanation C. Intra-site replication is configured to occur every 5 minutes by default. Replication will occur every 5 minutes to its replication partners and can have a maximum of 3 hops. It can take up to 15 minutes to replicate to other domain controllers.

Explanation D. Intra-site replication is configured to occur every 5 minutes by default.

Explanation E. Intra-site replication is configured to occur every 5 minutes by default.

5. Answer: D

Explanation A. There is no mention that separate security policies are required.

Explanation B. Sites can be created to control replication traffic over slow WAN connections.

Explanation C. The fact that the company already has an Internet presence will not impact the number of domains created.

Explanation D. The administrative model used by an organization will define whether centralized or decentralized administration is used. This will impact the number of domains created.

6. Answer: A

Explanation A. This would allow the organization to maintain centralized administration while allowing administrators within each distribution center to perform daily administrative tasks.

Explanation B. The organization wants to maintain centralized administrative control and reduce the number of domains.

Explanation C. The organization wants to maintain centralized administrative control and reduce the number of domains.

Explanation D. The administrative model used by an organization will define whether centralized or decentralized administration is used. This will impact the number of domains created.

7. Answer: D

Explanation A. Modifications to the schema are forest wide.

Explanation B. Changes to the schema are forest wide and will affect all domains within a single forest.

Explanation C. Changes to the schema are forest wide and will affect all domains within a single forest.

Explanation D. Changes to the schema are forest wide and will affect all domains within a single forest. Therefore, if MVP Kids is included in the same forest as MCP Inc., changes to the schema will affect the entire organization. To avoid this, a new forest should be created for MVP Kids.

8. Answer: D

Explanation A. Trusts are not automatically created between forest root domains.

Explanation B. Shortcut trusts are configured between domains in the same forest to shorten the trust path and improve the authentication process. Trusts are not automatically configured between forest root domains.

Explanation C. An NTLM trust would only be required if the infrastructure still needed to support Windows NT 4.0 domains.

Explanation D. By creating a forest trust, a two-way transitive trust is established between two forest root domains to establish a transitive trust relationship between all domains in each forest.

9. Answers: B, D, E, F

Explanation A. There can only be a single schema master per forest.

Explanation B. Placing a Global Catalog within each distribution office will ensure that users always have access to directory information. It will also decrease the amount of time it takes to search for network objects.

Explanation C. There is only one domain naming master within each forest.

Explanation D. Placing a domain controller within each location will ensure that users do not have to log on via a WAN link. This will decrease authentication time and increase availability.

Explanation E. This will ensure that users can still resolve host names in the event that a WAN link is unavailable.

Explanation F. WINS is required by pre-Windows 2000 clients. Placing a WINS server within each site will ensure users can resolve NetBIOS names in the event a WAN link is unavailable.

10. Answer: C

Explanation A. The existing model uses a single master account domain and four resource domains. Therefore, at least four one-way trusts must exist in which the resource domains trust the master account domain.

Explanation B. The existing model uses a single master account domain and four resource domains. Therefore, at least four one-way trusts must exist in which the resource domains trust the master account domain.

Explanation C. The existing model uses a single master account domain and four resource domains. Therefore, at least four one-way trusts must exist in which the resource domains trust the master account domain.

Explanation D. The existing model uses a single master account domain and four resource domains. Therefore, at least four one-way trusts must exist in which the resource domains trust the master account domain.

11. Answer: B

Explanation A. This would result in client computers on various subnets being assigned an incorrect IP address for their local default gateway.

Explanation B. By configuring the 003 router option at the scope level, each subnet can be assigned the correct IP address for their local default gateway.

Explanation C. Because all users in a specific subnet will require the same IP address of the default gateway, it is easier to configure this option at the scope level.

Explanation D. Because all users in a specific subnet will require the same IP address of the default gateway, it is easier to configure this option at the scope level.

12. Answer: C

Explanation A. HOSTS files must be manually updated; therefore, this would increase the administrative overhead.

Explanation B. A secondary DNS server must retrieve the zone file updates from a master name server. Therefore, this option would result in replication traffic on the WAN link.

Explanation C. Caching-only DNS servers are not responsible for zone information. They resolve queries and cache the results. Therefore, this would provide a DNS solution with minimal administration and no zone transfer traffic on the WAN link.

Explanation D. WINS is used to resolve NetBIOS names to IP addresses.

13. Answer: C

Explanation A. The infrastructure master is responsible for updating changes in each group.

Explanation B. The infrastructure master is responsible for updating changes in each group.

Explanation C. The infrastructure master is responsible for updating changes in each group.

Explanation D. The infrastructure master is responsible for updating changes in each group.

14. Answer: B

Explanation A. Because resource domains do not contain user accounts, they should be placed in OUs and the necessary users given authority over the appropriate OU.

Explanation B. Because resource domains do not contain user accounts, they should be placed in OUs and the necessary users given authority over the appropriate OU.

Explanation C. Within Active Directory, domains are the security boundaries.

Explanation D. Separate domains are not required for delegation of authority.

15. Answer: D

Explanation A. To seize the role from a domain controller that is no longer online, you must use Ntdsutil.

Explanation B. To seize the role from a domain controller that is no longer online, you must use Ntdsutil.

Explanation C. To transfer the role, both domain controllers must be online.

Explanation D. To seize the role from a domain controller that is no longer online, you must use Ntdsutil.

Case Study 2

1. Answer: C

Explanation A. The speed of WAN connections between locations will not impact the forest design. It will impact the site design.

Explanation B. A single Active Directory domain can contain millions of objects. Therefore, the number of users will not impact the forest design.

Explanation C. The administrative model will impact the forest design. For example, if complete autonomy is required, separate forests should be created.

Explanation D. The existing DNS structure will not impact the forest design.

2. Answers: A, D

Explanation A. By default, there is no authentication between forests. Trust relationships must be established.

Explanation B. Each forest will maintain its own Global Catalog Server.

Explanation C. Resources can be shared between forests if a trust relationship is established.

Explanation D. Because each forest will have its own schema, schema updates will have to be performed on the schema master in each forest.

3. Answer: A

Explanation A. Before you can create forest trusts, the forest functional level must be raised to Windows Server 2003.

Explanation B. Resources can be shared if a trust relationship exists between the forests.

Explanation C. The functional level of the domains has been raised because universal groups have been created. This feature is not available with the Windows 2000 functional level.

Explanation D. A user does not need to be a member of the Schema Admins group to create a trust relationship.

4. Answer: D

Explanation A. The site design should call for five sites, one for each location. Because each site spans multiple domains, a domain controller for each domain should be placed in each site. Therefore, a minimum of 10 domain controllers should be deployed.

Explanation B. The site design should call for five sites, one for each location. Because each site spans multiple domains, a domain controller for each domain should be placed in each site. Therefore, a minimum of 10 domain controllers should be deployed.

Explanation C. The site design should call for five sites, one for each location. Because each site spans multiple domains, a domain controller for each domain should be placed in each site. Therefore, a minimum of 10 domain controllers should be deployed.

Explanation D. The site design should call for five sites, one for each location. Because each site spans multiple domains, a domain controller for each domain should be placed in each site. Therefore, a minimum of 10 domain controllers should be deployed.

5. Answers: B, D

Explanation A. The company wants to implement separate internal and external namespaces.

Explanation B. This would provide an appropriate representation of the home furnishings division. The internal and external namespace remain separate.

Explanation C. The company wants to use separate internal and external namespaces. This name would be a subdomain of the company's current Internet domain.

Explanation D. This domain name would be an appropriate forest root name for the appliance division. The name is separate from the company's Internet name.

6. Answer: B

Explanation A. Resources must be shared between domains in each forest; therefore, a forest trust should be created. This creates a two-way transitive trust between the two root domains.

Explanation B. Resources must be shared between domains in each forest; therefore, a forest trust should be created. This creates a two-way transitive trust between the two root domains. Keep in mind that the forest trust itself is not transitive. This means if a third forest is created, the trust would not go through to it.

Explanation C. Resources must be shared between domains in each forest; therefore, a forest trust should be created. This creates a two-way transitive trust between the two root domains.

Explanation D. Resources must be shared between domains in each forest; therefore, a forest trust should be created. This creates a two-way transitive trust between the two root domains.

7. Answers: A, E

Explanation A. The number of locations will have a direct impact on the site design.

Explanation B. The operating system installed on client computers will not impact the site design. It will, however, impact how the clients interoperate with Active Directory.

Explanation C. The existing domain model will not impact the site design. It will impact the logical design of Active Directory.

Explanation D. The number of domain controllers in the existing infrastructure will not impact the site design.

Explanation E. The amount of bandwidth available for each WAN connection will impact the site design.

8. Answer: D

Explanation A. Although Active Directory integrated zones are more efficient and secure than standard zones, it will not lessen the impact of replication traffic.

Explanation B. GPOs configured at the OU level must still be replicated between domain controllers.

Explanation C. Domain local groups are required to assign permissions to resources.

Explanation D. RPC over IP site links allow you to configure a replication schedule. This can assist in minimizing the impact of replication traffic.

9. Answer: C

Explanation A. This would result in replication over WAN links occurring on demand.

Explanation B. The scenario states that the company currently has no plans to upgrade the existing WAN connections.

Explanation C. One site should be created for each location to optimize replication over WAN links. IP site links will allow you to schedule when replication should occur.

Explanation D. The scenario states that the company currently has no plans to upgrade the existing WAN connections.

10. Answer: D

Explanation A. Because the WAN connection between Vancouver and Calgary is slower than the WAN connection between Vancouver and Edmonton, the cost of the site link should be higher than the default value.

Explanation B. Because the WAN connection between Vancouver and Calgary is slower than the WAN connection between Vancouver and Edmonton, the cost of the site link should be higher than the default value.

Explanation C. Because the WAN connection between Vancouver and Calgary is slower than the WAN connection between Vancouver and Edmonton, the cost of the site link should be higher than the default value.

Explanation D. Because the WAN connection between Vancouver and Calgary is slower than the WAN connection between Vancouver and Edmonton, the cost of the site link should be higher than the default value.

11. Answer: B

Explanation A. You cannot configure a computer as a domain controller during the installation of Windows Server 2003. This must be done after installation.

Explanation B. Before configuring a server as a domain controller, you should always perform a backup of any data. A computer can be promoted to a domain controller using the DCPROMO command.

Explanation C. It is not necessary to reinstall the operating system.

Explanation D. All data should be backed up first. A computer is promoted to a domain controller using the DCPROMO command.

12. Answer: A

Explanation A. One site will exist by default. Additional sites must be created by an administrator.

Explanation B. One site will exist by default. Additional sites must be created by an administrator.

Explanation C. One site will exist by default. Additional sites must be created by an administrator.

Explanation D. One site will exist by default. Additional sites must be created by an administrator.

13. Answer: B

Explanation A. The site link name is used to easily identify the site link.

Explanation B. Because the WAN connection from Vancouver to Calgary is the slower connection, the cost of the link should be increased.

Explanation C. RPC over IP is required to schedule replication.

Explanation D. The interval defines how often replication can occur over the site link. The default value is 180.

14. Answers: A, C

Explanation A. Windows NT Server 4.0 can be upgraded directly to Windows Server 2003 Standard Edition.

Explanation B. Windows 2000 Professional cannot be directly upgraded to Windows Server 2003 Standard Edition.

Explanation C. Windows 2000 Server can be upgraded directly to Windows Server 2003 Standard Edition.

Explanation D. Windows NT Server 4.0 Enterprise Edition cannot be directly upgraded to Windows Server 2003 Standard Edition.

15. Answer: B

Explanation A. ICS would be a more effective solution because it is simpler to configure and administer.

Explanation B. ICS would allow all computers to share a single Internet connection. It is a cost-effective solution and simple to administer.

Explanation C. ICS would be a more effective solution because it is simpler to configure and administer.

Explanation D. IAS provides centralized authentication and accounting for a remote access solution.