Module 13

Managing Sites and Active Directory® Replication

Contents:

Lesson 1: Configure Sites and Subnets 13-3

Lab A: Configure Sites and Subnets 13-17

Lesson 2: Configure Replication 13-21

Lab B: Configure Replication 13-39

Module Overview

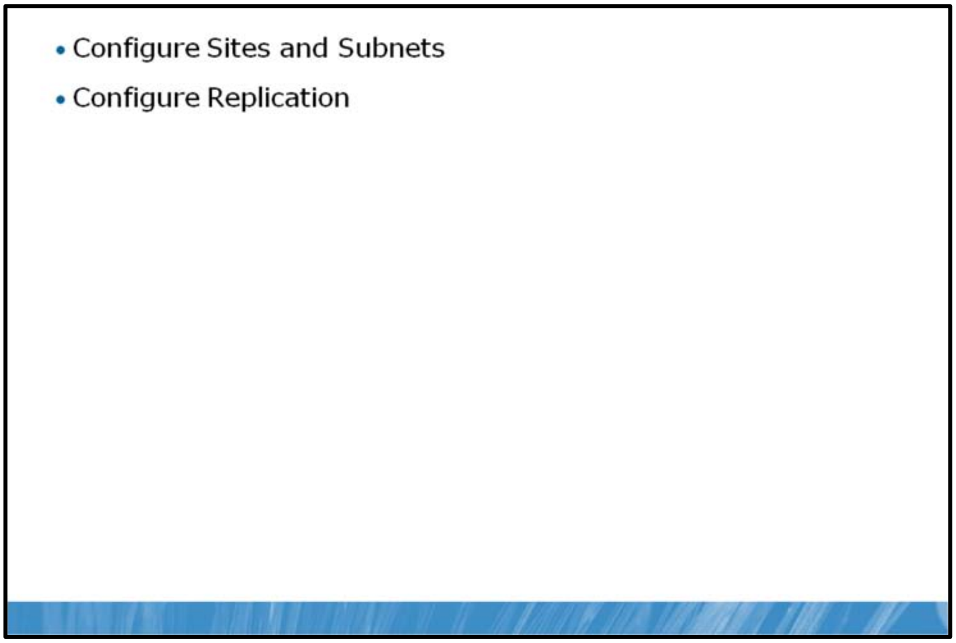

- Configure Sites and Subnets
- Configure Replication

You learned in previous modules that domain controllers in a Windows Server® 2008 or Windows Server 2008 R2 domain are peers. Each domain controller maintains a copy of the Active Directory® Domain Services (AD DS) and performs similar services to support authentication of security principals, and changes made on any one domain controller are replicated to all other domain controllers. As an administrator of a Windows® enterprise, you need to ensure that efficient authentication is provided and that replication between domain controllers is optimized. Active Directory sites are the core components of the directory service that supports the goals of service localization and replication. In this module, you will learn how to create a distributed directory service that supports domain controllers in portions of your network that are separated by expensive, slow, or unreliable links. You'll learn where domain controllers should be placed and how to manage replication and service utilization. You'll also learn how to control which data is replicated to each domain controller by configuring global catalogs and application partitions.

Objectives

After completing this module, you will be able to:

- Configure sites and subnets.
- Configure global catalog servers and application directory partitions.
- Configure replication topology with connection objects, bridgehead servers, site links, and site link bridges.

Lesson 1
Configure Sites and Subnets

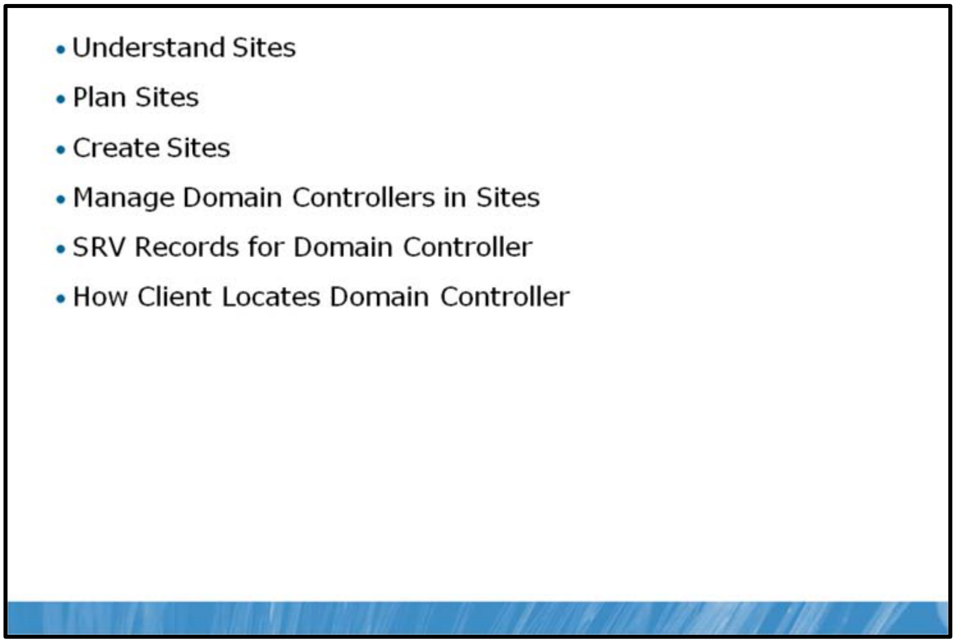

- Understand Sites
- Plan Sites
- Create Sites
- Manage Domain Controllers in Sites
- SRV Records for Domain Controller
- How Client Locates Domain Controller

Active Directory represents users as user objects in the directory service. It represents machines as computer objects. It represents network topology with objects as sites and subnets. Active Directory site objects are used to manage replication and service localization and, fortunately, in many environments, the configuration of sites and subnets can be quite straightforward. In this lesson, you will learn the fundamental concepts and techniques required to configure and manage sites and subnets.

Objectives

After completing this lesson, you will be able to:

- Identify the role of sites and subnets.
- Describe the process with which a client locates a domain controller.
- Configure sites and subnets.
- Manage domain controller server objects in sites.

Understand Sites

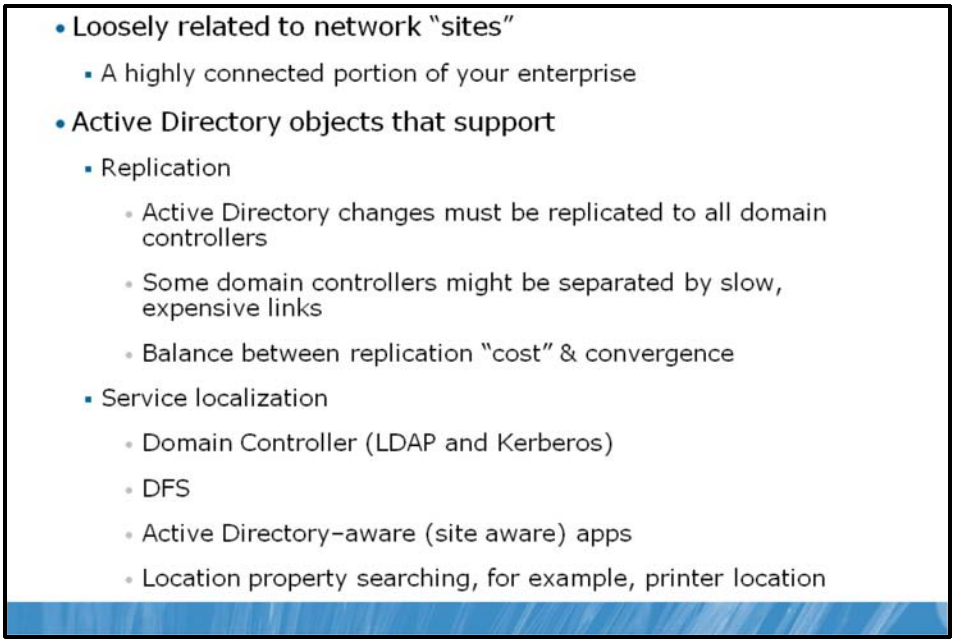

- Loosely related to network "sites"
  - A highly connected portion of your enterprise
- Active Directory objects that support
  - Replication
    - Active Directory changes must be replicated to all domain controllers
    - Some domain controllers might be separated by slow, expensive links
    - Balance between replication "cost" & convergence
  - Service localization
    - Domain Controller (LDAP and Kerberos)
    - DFS
    - Active Directory–aware (site aware) apps
    - Location property searching, for example, printer location

When administrators describe their network infrastructure, they often mention how many sites comprise their enterprise. To most administrators, a site is a physical location, an office, or a city. Sites are connected by links—network links that might be as basic as dial-up connections or as sophisticated as fiber links. Together, the physical locations and links make up the network infrastructure.

Active Directory represents the network infrastructure with objects called sites and site links, and though the words are similar, the objects are not identical to the sites and links described by administrators. This lesson focuses on sites, and Lesson 2 discusses site links.

It's important to understand the properties and roles of sites in Active Directory to understand the distinction between Active Directory sites and network sites. Active Directory sites are objects in the directory, specifically in the Configuration container (CN=Configuration, DC=*forest root domain*). These objects are used to achieve two service management tasks:

- To manage replication traffic
- To facilitate service localization

Replication Traffic

Replication is the transfer of changes between domain controllers. When you add a user or change a user's password, the change you make is committed to the directory by one domain controller. That change must be communicated to all other domain controllers in the domain.

Active Directory assumes that there are two types of networks within your enterprise: highly connected and less highly connected. Conceptually, a change made to Active Directory should replicate immediately to other domain controllers within the highly connected network in which the change was made. However, you might not want the change to replicate immediately over a slower, more expensive, or less reliable link to another site. Instead, you might want to manage replication over less highly connected segments of your enterprise to optimize performance, reduce costs, or manage bandwidth.

An Active Directory site represents a highly connected portion of your enterprise. When you define a site, the domain controllers within the site replicate changes almost instantly. Replication between sites can be scheduled and managed.

Service Localization

Active Directory is a distributed service. Assuming you have at least two domain controllers, multiple servers or domain controllers provide the same services of authentication and directory access. If you have more than one network site, and if you place a domain controller in each, you can encourage clients to authenticate against the domain controller in their site. This is an example of service localization.

Active Directory sites help you localize services, including those provided by domain controllers. During logon, Windows clients are automatically directed to a domain controller in their site. If a domain controller is not available in their site, they are directed to a domain controller in another, closest site that will be able to authenticate the client efficiently.

Plan Sites

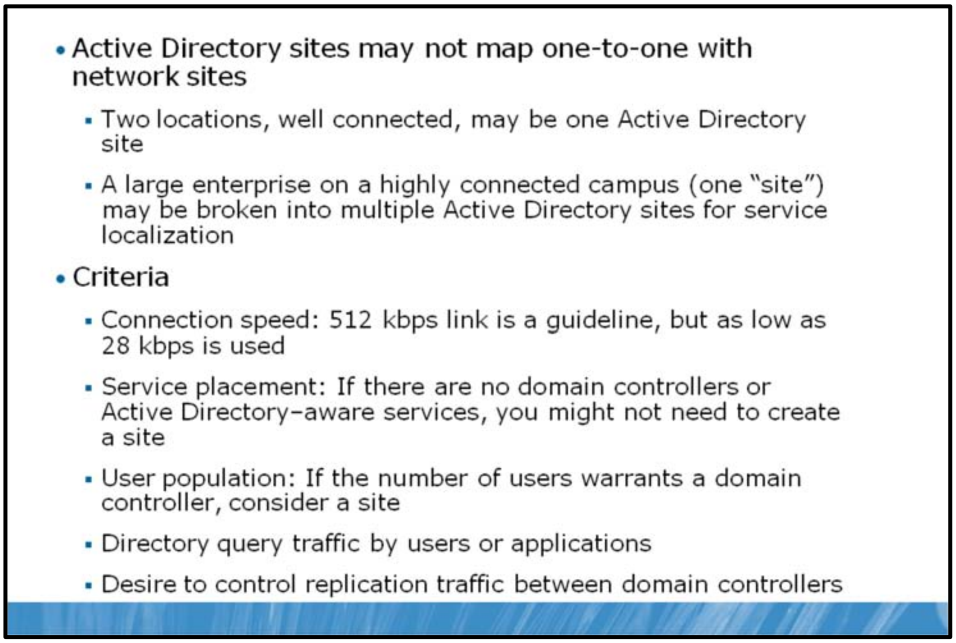

- Active Directory sites may not map one-to-one with network sites
  - Two locations, well connected, may be one Active Directory site
  - A large enterprise on a highly connected campus (one "site") may be broken into multiple Active Directory sites for service localization
- Criteria
  - Connection speed: 512 kbps link is a guideline, but as low as 28 kbps is used
  - Service placement: If there are no domain controllers or Active Directory–aware services, you might not need to create a site
  - User population: If the number of users warrants a domain controller, consider a site
  - Directory query traffic by users or applications
  - Desire to control replication traffic between domain controllers

Because sites are used to optimize replication and enable service localization, you must spend time in designing the Active Directory site structure. Active Directory sites might not map one to one with your network's sites. Consider the following two scenarios:

- You have offices in two distinct locations. You place one domain controller in each location. The locations are highly connected. Then, to improve performance, you decide to configure a single Active Directory site, which includes both locations.
- You have an enterprise on a large, highly connected campus. From a replication perspective, the enterprise could be considered a single site. However, you want to encourage clients to use distributed services in their location, so you configure multiple sites to support service localization.

Therefore, an Active Directory site can include more than one network site or be a subset of a single network site. The key is to remember that sites serve both replication management and service localization roles. Several characteristics of your enterprise can be used to help you determine which sites are necessary.

Connection Speed

An Active Directory site represents a unit of the network that is characterized by fast, reliable, inexpensive connectivity. Documentation suggests that the slowest link speed within a site should be no less than 512 kilobits per second (Kbps). However, this guidance is not immutable. Some organizations have links as slow as 56 Kbps or even 28 kbps within a site.

Service Placement

Because Active Directory sites manage Active Directory replication and service localization, it is not useful to create a site for a network location that does not host a domain controller or other Active Directory–aware service such as a replicated DFS resource.

Note Domain controllers are only one distributed service in a Windows enterprise. Other services, such as replicated DFS resources, are site-aware as well. You might configure sites to localize services other than authentication, in which case, you will have sites without domain controllers.

User Population

Concentration of users can also influence your site design, though indirectly. If a network location has a sufficient number of users for whom the inability to authenticate would be problematic, place a domain controller in the location to support authentication within the location. After a domain controller or other distributed service is placed in the location to support those users, you might want to manage Active Directory replication to the location or localize service use by configuring an Active Directory site to represent the location.

Summarizing Site Planning Criteria

Every Active Directory forest includes at least one site. When you instantiate a forest with the first domain controller, the default site created is named Default-First-Site-Name. You should create additional sites when:

- A part of the network is separated by a slow link.
- A part of the network has enough users to warrant hosting domain controllers or other services in that location.
- Directory query traffic warrants a local domain controller.
- You want to control service localization.
- You want to control replication between domain controllers.

A Note on Server (Domain Controller) Placement

Network administrators often want to know when placing a domain controller in a remote site is recommended. The answer depends on the resources required by users in the site and the tolerance for downtime. For example, if users in a remote site perform all tasks by accessing resources in the data center, and if the link to the remote site fails, the users cannot access the resources they require, and a local domain controller would not improve the situation. However, if users access resources in the remote site and the link fails, a local domain controller can continue to provide authentication for users, and they can continue to work with their local resources.

In most branch office scenarios, there are resources in the branch office that users require to perform their work-related tasks. Those resources, if not stored on the user's own computer, require domain authentication of the user. Therefore, a domain controller is generally recommended. The introduction of read-only domain controllers (RODCs) in Windows Server 2008 reduces the risk and management burden of domain controllers in branch offices, so it will be easier for most organizations to deploy domain controllers in each network location.

Create Sites

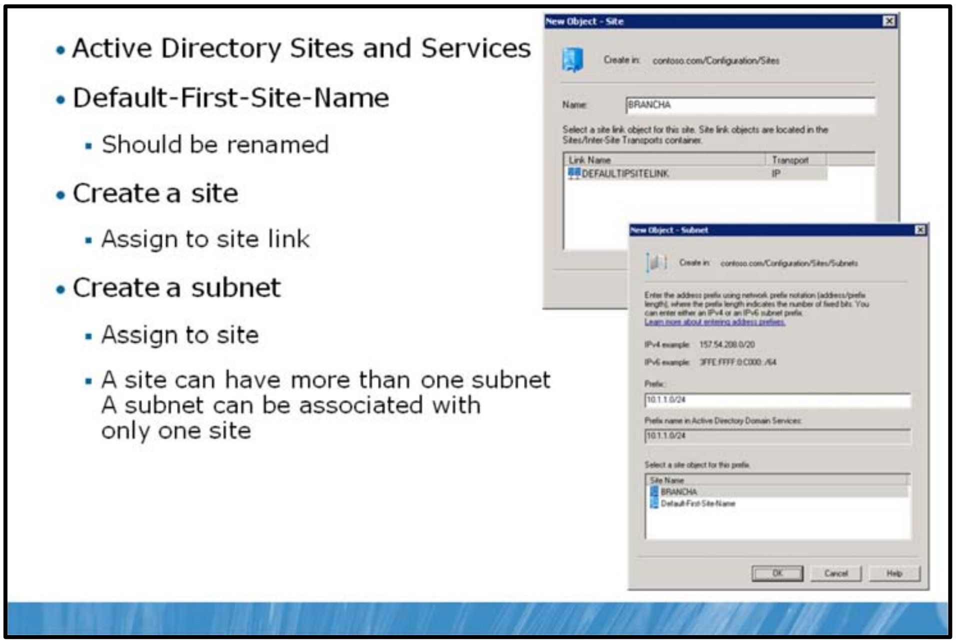

- Active Directory Sites and Services
- Default-First-Site-Name
  - Should be renamed
- Create a site
  - Assign to site link
- Create a subnet
  - Assign to site
  - A site can have more than one subnet
    A subnet can be associated with only one site

New Object - Site

Create in: contoso.com/Configuration/Sites

Name: BRANCHA

Select a site link object for this site. Site link objects are located in the Sites/Inter-Site Transports container.

Link Name	Transport
DEFAULTIPSITELINK	IP

New Object - Subnet

Create in: contoso.com/Configuration/Sites/Subnets

Enter the address prefix using network prefix notation (address/prefix length), where the prefix length indicates the number of fixed bits. You can enter either an IPv4 or an IPv6 subnet prefix.
Learn more about entering address prefixes.

IPv4 example: 157.54.208.0/20

IPv6 example: 3FFE:FFFF:0:C000::/64

Prefix:
10.1.1.0/24

Prefix name in Active Directory Domain Services:
10.1.1.0/24

Select a site object for this prefix.

Site Name
BRANCHA
Default-First-Site-Name

Sites and replication are managed by using the Active Directory Sites and Services snap-in. To define an Active Directory site, you will create an object of class site. The site object is a container that manages replication for domain controllers in the site. You will also create one or more subnet objects. A subnet object defines a range of IP addresses and is linked to one site. Service localization is attained when a client's IP address can be associated with a site through the relationship between the subnet object and the site object.

To create a site:

1. Right-click the **Sites** node in **Active Directory Sites and Services,** and then click **New Site**.
2. In the **New Object – Site** dialog box that appears, type a site name, and select a site link.

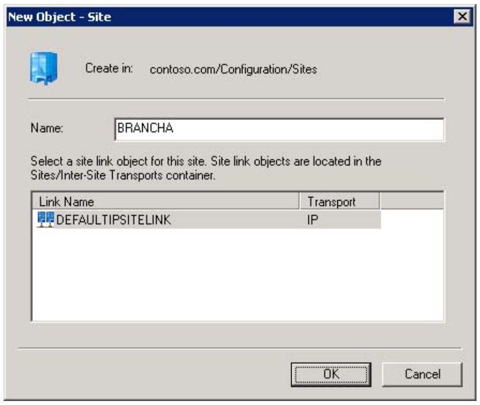

New Object - Site

Create in: contoso.com/Configuration/Sites

Name: BRANCHA

Select a site link object for this site. Site link objects are located in the Sites/Inter-Site Transports container.

Link Name	Transport
DEFAULTIPSITELINK	IP

The default site link, DEFAULTIPSITELINK, will be the only site link available to you until you create additional site links, as discussed in Lesson 2.

After creating a site, you can right-click the site and select **Rename** to rename it. Rename the Default-First-Site-Name site to reflect a site name that is aligned with your business and network topology.

Sites are useful only when a client or server knows the site to which it belongs. This is typically achieved by associating the system's IP address with a site, and subnet objects achieve this association.

To create a subnet object:

3. In the **Active Directory Sites and Services** snap-in, right-click the **Subnets** node and click **New Subnet**. The **New Object – Subnet** dialog box appears.
4. Enter the network prefix and subnet mask length.

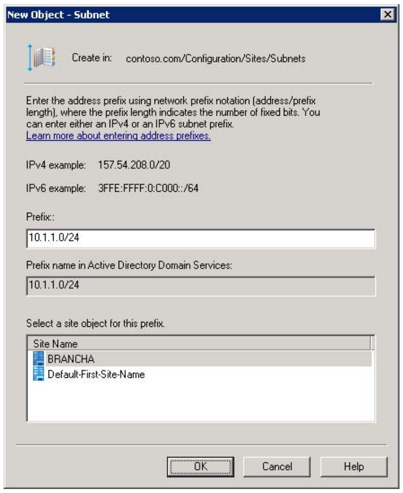

New Object - Subnet

Create in: contoso.com/Configuration/Sites/Subnets

Enter the address prefix using network prefix notation (address/prefix length), where the prefix length indicates the number of fixed bits. You can enter either an IPv4 or an IPv6 subnet prefix.
Learn more about entering address prefixes.

IPv4 example: 157.54.208.0/20

IPv6 example: 3FFE:FFFF:0:C000::/64

Prefix:: 10.1.1.0/24

Prefix name in Active Directory Domain Services: 10.1.1.0/24

Select a site object for this prefix.

Site Name
BRANCHA
Default-First-Site-Name

The subnet object is defined as a range of addresses that use network prefix notation. For example, to enter a subnet representing the addresses 10.1.1.1 to 10.1.1.254 with a 24-bit subnet mask, the prefix would be 10.1.1.0/24. For more information about entering addresses, click the Learn More About Entering Address Prefixes link in the New Object – Subnet dialog box.

5. After entering the network prefix, select the site object with which the subnet is associated.

   A subnet can be associated with only one site; however, a site can have more than one subnet linked to it. The **Properties** dialog box of a site, shown in the following screen shot, shows the subnets associated with the site. You cannot change the subnets in this dialog box. Instead, you must open the properties of the subnet, shown in the following screen shot, to change the site to which the subnet is linked.

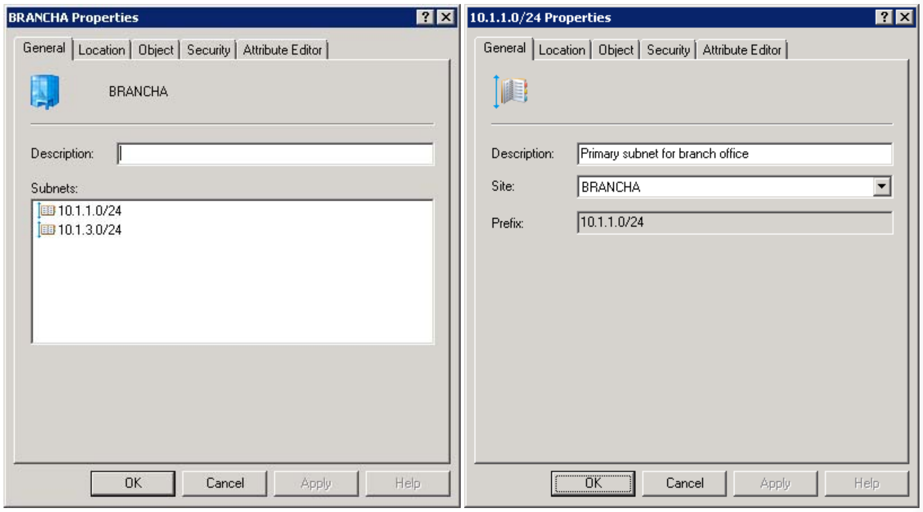

BRANCHA

Description:

Subnets:

10.1.1.0/24

10.1.3.0/24

Description: Primary subnet for branch office

Site: BRANCHA

Prefix: 10.1.1.0/24

Note In your production environment, define each IP subnet as an Active Directory subnet object. If a client's IP address is not included within a subnet range, the client is unable to determine which Active Directory site it belongs to and which can lead to performance and functionality problems. Don't forget backbone subnets and subnets used for remote access such as virtual private network (VPN) address ranges.

Manage Domain Controllers in Sites

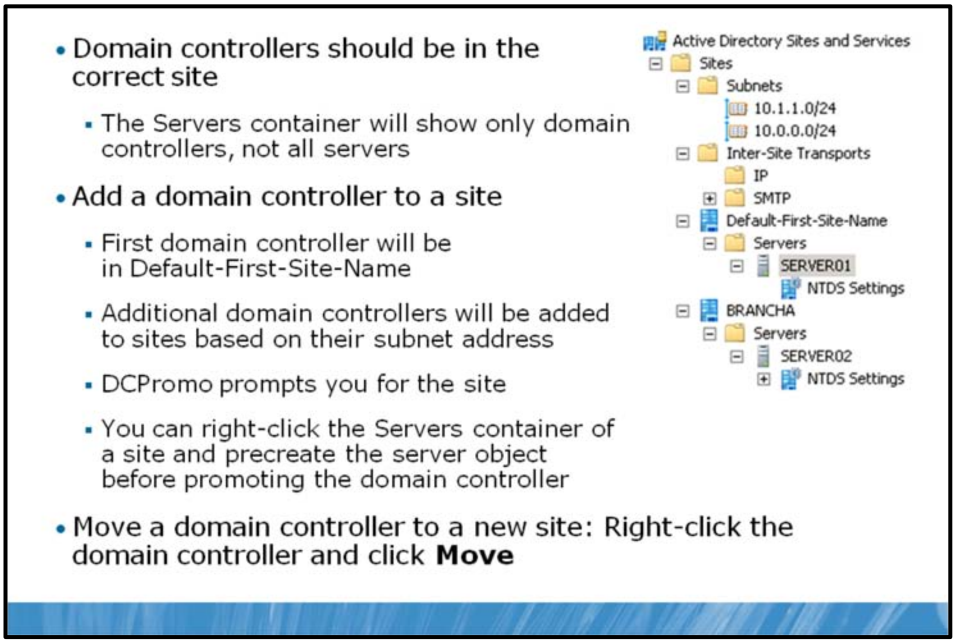

- Domain controllers should be in the correct site
  - The Servers container will show only domain controllers, not all servers
- Add a domain controller to a site
  - First domain controller will be in Default-First-Site-Name
  - Additional domain controllers will be added to sites based on their subnet address
  - DCPromo prompts you for the site
  - You can right-click the Servers container of a site and precreate the server object before promoting the domain controller
- Move a domain controller to a new site: Right-click the domain controller and click **Move**

Sometimes, you might need to manage domain controllers in Active Directory sites. Some such scenarios include:

- You create a new site and move an existing domain controller to it.
- You demote a domain controller.
- You promote a new domain controller.

When you create your Active Directory forest, the first domain controller is automatically placed under the site object named Default-First-Site-Name. You can see the domain controller SERVER01.contoso.com in the following screen shot.

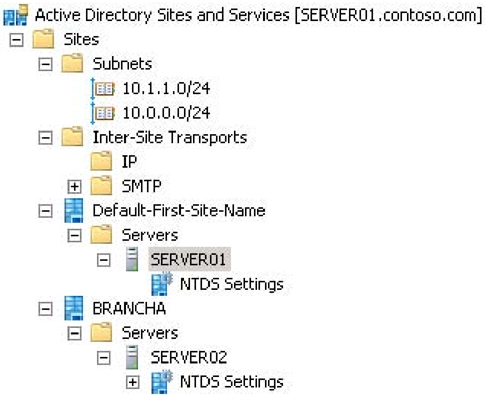

Additional domain controllers will be added to sites based on their IP addresses. For example, if a server with IP address 10.1.1.17 is promoted to a domain controller, the server will automatically be added to the BRANCHA site because the 10.1.1.0/24 subnet was associated with the BRANCHA site (see the previous slide). The previous screen shot shows SERVER02 in the BRANCHA site.

Each site contains a Servers container, which itself contains an object for each domain controller in the site. The Servers container in a site should show only domain controllers; not all servers. When you promote a new domain controller, it is placed in the site associated with its IP address, by default. However, the Active Directory Domain Services Installation Wizard will enable you to specify another site. You can also precreate the server object for the domain controller in the correct site by right-clicking the Servers container in the appropriate site and choosing a Server from the New menu.

Finally, you can move the domain controller to the correct site after installation by right-clicking the server and choosing **Move**. In the **Move Server** dialog box, select the new site and click **OK**. The domain controller is moved. It is a best practice to place a domain controller in the site object that is associated with the domain controller's IP address. If a domain controller is multihomed, it can belong to only one site. If a site has no domain controllers, users will still be able to log on to the domain; their logon requests will be handled by a domain controller in an adjacent site or another domain controller in the domain.

SRV Records for Domain Controller

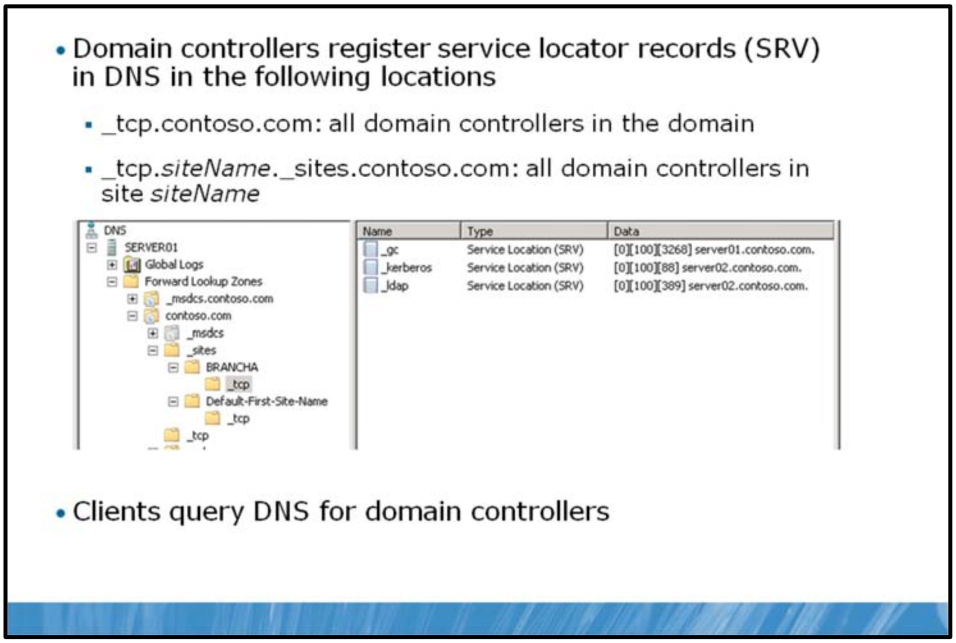

- Domain controllers register service locator records (SRV) in DNS in the following locations
  - _tcp.contoso.com: all domain controllers in the domain
  - _tcp.*siteName*._sites.contoso.com: all domain controllers in site *siteName*

Name	Type	Data
_gc	Service Location (SRV)	[0][100][3268] server01.contoso.com.
_kerberos	Service Location (SRV)	[0][100][88] server02.contoso.com.
_ldap	Service Location (SRV)	[0][100][389] server02.contoso.com.

- Clients query DNS for domain controllers

When a domain controller is added to the domain, it advertises its services by creating Service Locator (SRV) records, also called locator records, in Domain Name System (DNS). Unlike host records (A records), which map host names to IP addresses, SRV records map services to host names. The domain controller advertises its ability to provide authentication and directory access by registering Kerberos and Lightweight Directory Access Protocol (LDAP) SRV records. These SRV records are added to several folders within the DNS zones for the forest. The first folder is within the domain zone. It is called tcp, and it contains the SRV records for all domain controllers in the domain. The second folder is specific to this site, in which the domain controller is located, with the path _sites*sitename*_tcp, where *sitename* is the name of the site.

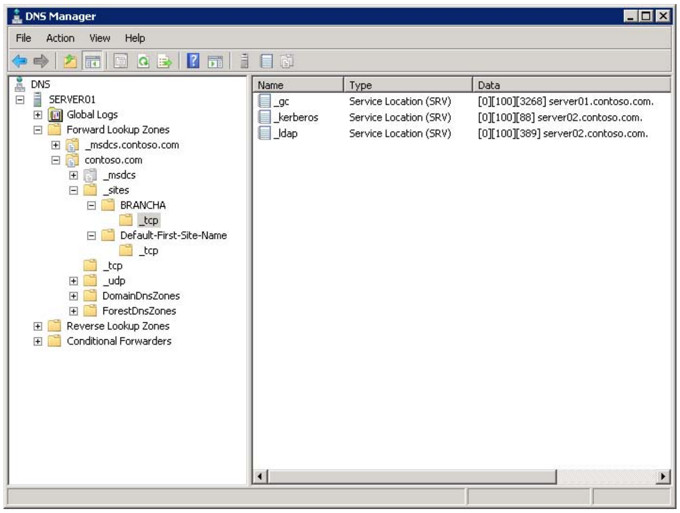

Name	Type	Data
_gc	Service Location (SRV)	[0][100][3268] server01.contoso.com.
_kerberos	Service Location (SRV)	[0][100][88] server02.contoso.com.
_ldap	Service Location (SRV)	[0][100][389] server02.contoso.com.

In the previous screen shot, you can see the Kerberos and LDAP SRV records for SERVER02.contoso.com in its site, sites\BRANCHA_tcp. You can also see the _tcp folder at the first level beneath the zone.

The same records are registered in several places in the msdcs.domainName zone, for example, _msdcs.contoso.com in the previous screen shot. This zone contains records for Microsoft Domain Controller Services. The underscore is a requirement of RFC 2052.

Locator records contain:

- **The service name and port.** This portion of the SRV record indicates a service with a fixed port. It does not have to be a well-known port. SRV records in Windows Server 2008 include LDAP (port 389), Kerberos (port 88), Kerberos Password protocol (KPASSWD, port 464), and global catalog services (port 3268).
- **Protocol.** Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) will be indicated as a transport protocol for the service. The same service can use both protocols, in separate SRV records. Kerberos records, for example, are registered for both TCP and UDP. Microsoft clients use only TCP, but UNIX clients can use TCP.
- **Host name.** The name corresponds to the A (Host) record for the server hosting the service. When a client queries for a service, the DNS server returns the SRV record and associated A records, so the client does not need to submit a separate query to resolve the IP address of a service.

The service name in the SRV record follows the standard DNS hierarchy, with components separated by dots. For example, the Kerberos service of a domain controller is registered as:

kerberos._tcp.*siteName***._sites.***domainName*

Reading this SRV record name right to left like other DNS records, it translates to:

- *domainName*: The domain or zone, for example contoso.com
- **_sites**: All sites registered with DNS
- *siteName*: The site of the domain controller registering the service
- **_tcp**: Any TCP-based services in the site
- **kerberos**: A Kerberos Key Distribution Center (KDC)that uses TCP as its transport protocol

How Client Locates Domain Controller

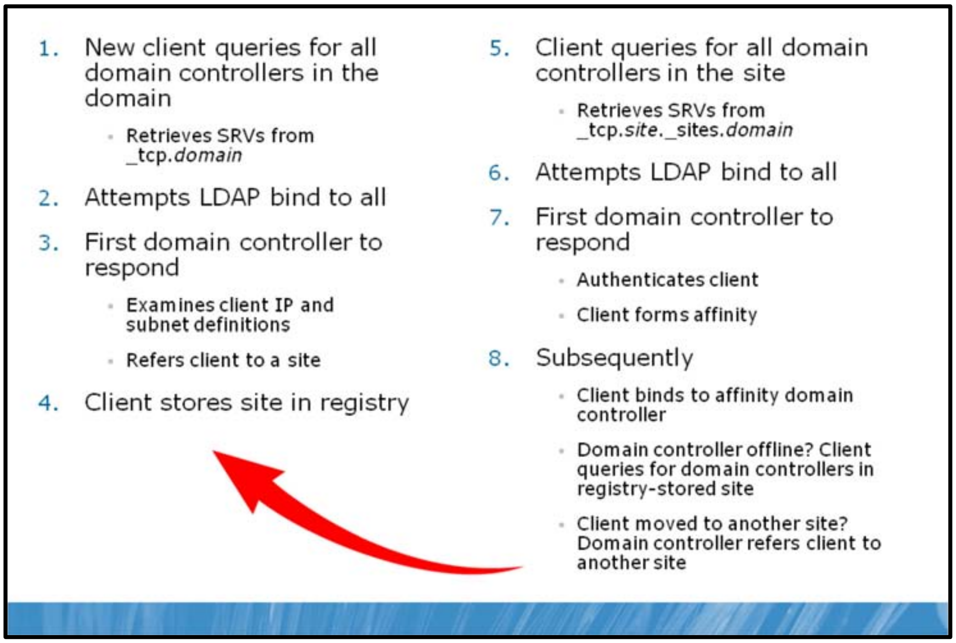

1. New client queries for all domain controllers in the domain
   - Retrieves SRVs from _tcp.*domain*
2. Attempts LDAP bind to all
3. First domain controller to respond
   - Examines client IP and subnet definitions
   - Refers client to a site
4. Client stores site in registry
5. Client queries for all domain controllers in the site
   - Retrieves SRVs from _tcp.*site*._sites.*domain*
6. Attempts LDAP bind to all
7. First domain controller to respond
   - Authenticates client
   - Client forms affinity
8. Subsequently
   - Client binds to affinity domain controller
   - Domain controller offline? Client queries for domain controllers in registry-stored site
   - Client moved to another site? Domain controller refers client to another site

Imagine a Windows client has just been joined to the domain. It restarts, receives an IP address from a Dynamic Host Control Protocol (DHCP) server, and is ready to authenticate to the domain. However, the client does not know where to find a domain controller. Therefore, the client queries the domain for a domain controller by querying the _tcp folder, which contains the SRV records for all domain controllers in the domain. DNS returns a list of all matching domain controllers and the client attempts to contact all of them on its first startup. The first domain controller that responds to the client examines the client's IP address, cross-references that address with subnet objects, and informs the client of the site to which the client belongs. The client stores the site name in its registry, and then queries for domain controllers in the site-specific tcp folder. DNS returns a list of all domain controllers in the site. The client attempts to bind with all, and the domain controller that responds first authenticates the client.

The client forms an affinity for this domain controller and attempts to authenticate with the same domain controller in the future. If the domain controller is unavailable, the client queries the site's _tcp folder again and attempts to bind with all domain controllers in the site. But what happens if the client is a mobile computer, such as a laptop? Imagine that the computer has been authenticating in the BRANCHA site and then the user brings the computer to the BRANCHB site. When the computer starts, it actually attempts to authenticate with its preferred domain controller into the BRANCHA site. The domain controller that notices the client's IP address is associated with BRANCHB and informs the client of its new site. Then, the client queries DNS for domain controllers in BRANCHB.

You can see how by storing subnet and site information in Active Directory and by registering services in DNS, a client is encouraged to use services in its site. This is the definition of service localization.

Site Coverage

What happens if a site has no domain controller? Sites can be used to direct users to local copies of replicated resources such as shared folders replicated within a Distributed File System (DFS) namespace, so you might have sites without a domain controller. In this case, a nearby domain controller will register its SRV records in the site in a process called site coverage. To be precise, a site without a domain controller will generally be covered by a domain controller in a site with the lowest cost to the site requiring coverage. You'll learn more about site link costs in Lesson 2. You can also manually configure site coverage and SRV record priority if you want to control authentication in sites without domain controllers. The URL just listed contains details about the algorithm that determines which domain controller automatically covers a site without a domain controller.

Lab A: Configure Sites and Subnets

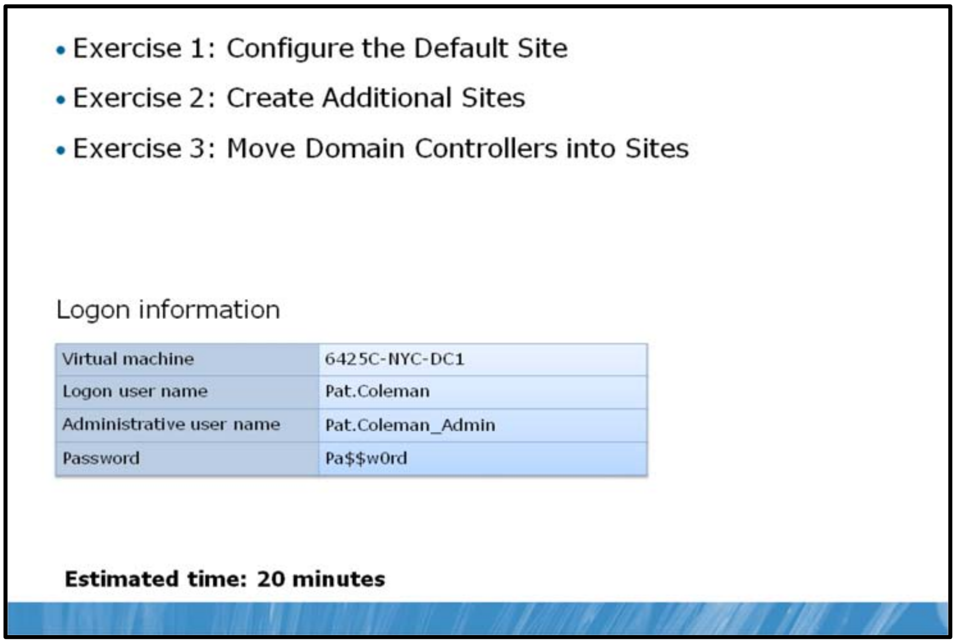

- Exercise 1: Configure the Default Site
- Exercise 2: Create Additional Sites
- Exercise 3: Move Domain Controllers into Sites

Logon information

Virtual machine	6425C-NYC-DC1
Logon user name	Pat.Coleman
Administrative user name	Pat.Coleman_Admin
Password	Pa$$w0rd

Estimated time: 20 minutes

Lab Setup

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V™ Manager, click **6425C-NYC-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
   - User name: **Pat.Coleman**
   - Password: **Pa$$w0rd**
   - Domain: **Contoso**

Lab Scenario

You are an administrator at Contoso, Ltd. You are preparing to improve the service localization and Active Directory replication of your enterprise. The previous administrator made no changes to the out-of-box configuration of sites and subnets. You want to begin the process of defining your physical topology in Active Directory.

Exercise 1: Configure the Default Site

In this exercise, you will rename the Default-First-Site-Name site and associate two subnets with the site.

The main tasks for this exercise are as follows:

1. Rename Default-First-Site-Name.
2. Create a subnet and associate it with a site.

▶ Task 1: Rename Default-First-Site-Name.

1. On NYC-DC1, run **Active Directory Sites and Services** with the user name **Pat.Coleman_Admin** and the password **Pa$$w0rd**.
2. Rename **Default-First-Site-Name** to **HEADQUARTERS**.

▶ Task 2: Create a subnet and associate it with a site.

- Create two subnets: **10.0.0.0/24** and **10.0.1.0/24**, and associate each with the **HEADQUARTERS** site. For the 10.0.0.0/24 subnet, provide a description that states "Server and back-end subnet. For the 10.0.1.0/24 subnet, provide a description that states "Client subnet".

Results: In this exercise, you named a site HEADQUARTERS and added two subnets (10.0.0.0/24 and 10.0.1.0/24) associated with the site.

Exercise 2: Create Additional Sites

In this exercise, you will create a second site and associate a subnet with it.

The main tasks for this exercise are as follows:

1. Create additional sites.
2. Create subnets and associate them with sites.

▶ Task 1: Create additional sites.

1. Create a site named **HQ-BUILDING-2**.
2. Create a site named **BRANCHA.**

   Connect both sites to the DEFAULTIPSITELINK site link object.

▶ Task 2: Create subnets and associate them with sites.

1. Create a subnet, **10.1.0.0/24**, and associate it with the **HQ-BUILDING-2** site. Provide a description for the subnet that states "Headquarters Building 2".
2. Create a subnet, **10.2.0.0/24**, and associate it with the **BRANCHA** site. Provide a description for the subnet that states "Branch Office A".

Results: In this exercise, you created two new sites, HQ-BUILDING-2 and BRANCHA, and associated them with the 10.1.0.0/24 and 10.2.0.0/24 subnets.

Exercise 3: Move Domain Controllers into Sites

In this exercise, you will move the NYC-DC2 and BRANCHDC02 domain controllers into the HQ-BUILDING-2 and BRANCHA sites.

The main task for this exercise is as follows:

- Move domain controllers to new sites.

▶ Task: Move domain controllers to new sites.

1. Move NYC-DC2 to the **HQ-BUILDING-2** site.
2. Move BRANCHDC02 to the **BRANCHA** site.

Results: In this exercise, you moved the NYC-DC2 and BRANCHDC02 domain controllers into the HQ-BUILDING-2 and BRANCHA sites.

Important Do not shut down the virtual machines after you finish this lab because the settings you have configured here will be used in subsequent labs in this module.

Lab Review Questions

Question: You have a site with 50 subnets, each with a subnet address of 10.0.x.0/24, and you have no other 10.0.x.0 subnets. What should you do to make it easier to identify the 50 subnets and associate them with a site?

Lesson 2
Configure Replication

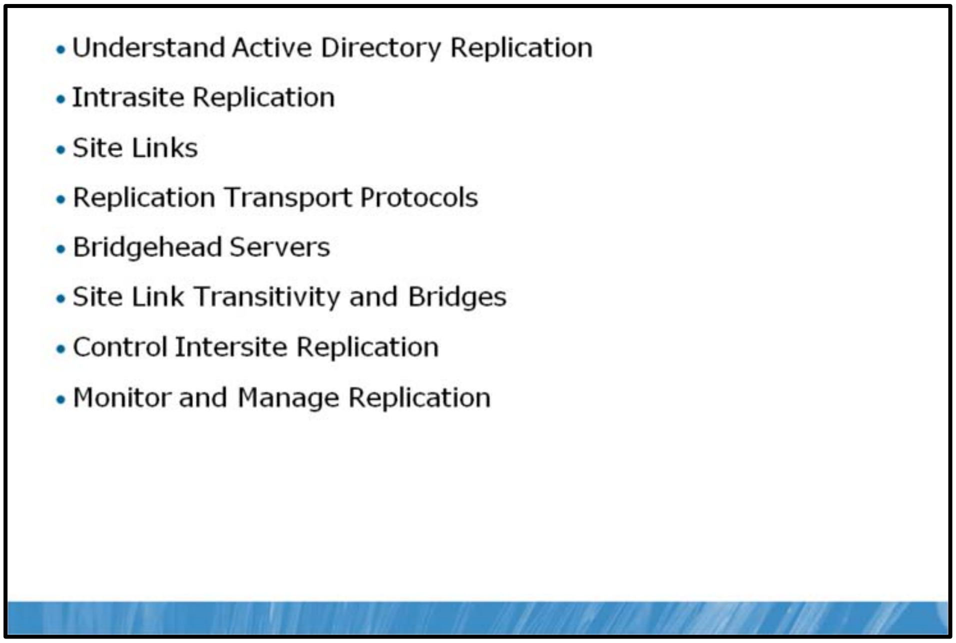

- Understand Active Directory Replication
- Intrasite Replication
- Site Links
- Replication Transport Protocols
- Bridgehead Servers
- Site Link Transitivity and Bridges
- Control Intersite Replication
- Monitor and Manage Replication

In Lesson 1, you learned how to create site and subnet objects that enable Active Directory and its clients to localize authentication and directory access. You also decided where domain controllers should be placed. In Lesson 2, you configured global catalog servers and application directory partitions, and you managed replication between domain controllers. In this lesson, you will learn how and when replication occurs. You'll discover why the default configuration of Active Directory supports effective replication and why you might modify that configuration so that replication is equally effective but more efficient based on your network topology.

Objectives

After completing this lesson, you will be able to:

- Create connection objects to configure replication between two domain controllers.
- Implement site links and site link costs to manage replication between sites.
- Designate preferred bridgehead servers.
- Understand notification and polling.
- Report and analyze replication with repadmin.exe.
- Perform Active Directory replication health checks with dcdiag.exe.

Understand Active Directory Replication

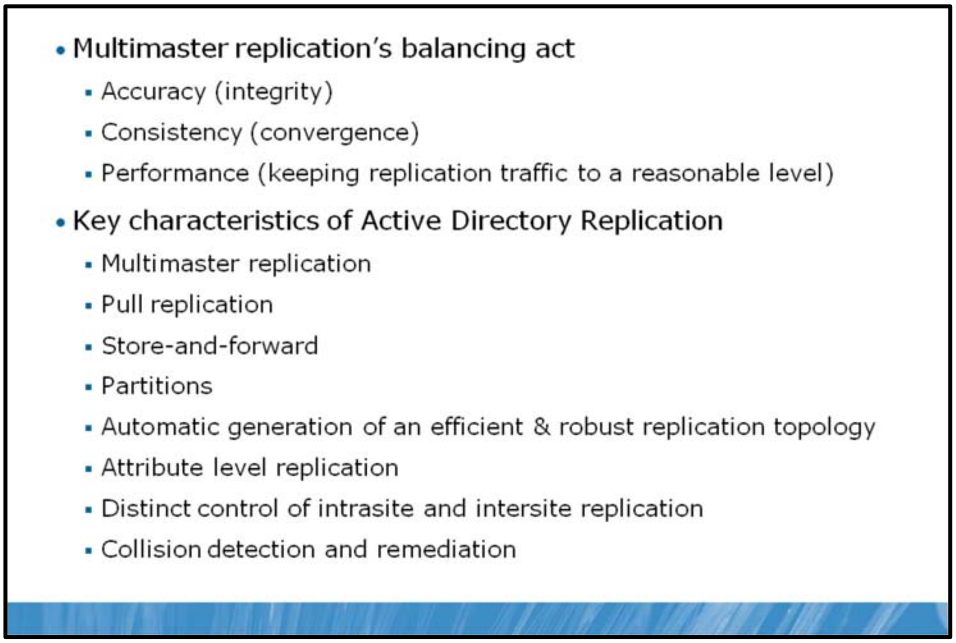

- Multimaster replication's balancing act
  - Accuracy (integrity)
  - Consistency (convergence)
  - Performance (keeping replication traffic to a reasonable level)
- Key characteristics of Active Directory Replication
  - Multimaster replication
  - Pull replication
  - Store-and-forward
  - Partitions
  - Automatic generation of an efficient & robust replication topology
  - Attribute level replication
  - Distinct control of intrasite and intersite replication
  - Collision detection and remediation

In previous lessons, you learned how to place domain controllers in network locations and how to represent those locations with site and subnet objects. You also learned about the replication of directory partitions (schema, configuration, and domain), the partial attribute set (global catalog), and application partitions. As you learn about Active Directory replication, the most important thing to remember is that it is designed so that, in the end, each replica on a domain controller is consistent with the replicas of that partition hosted on other domain controllers. Not all domain controllers will have exactly the same information in their replicas at any one moment in time because changes are constantly being made to the directory. However, Active Directory replication ensures that all changes to a partition are transferred to all replicas of the partition. Active Directory replication balances accuracy (or integrity) and consistency (called *convergence*) with performance (keeping replication traffic to a reasonable level).

The key characteristics of Active Directory replication are:

- **Multimaster replication.** Any domain controller except RODCs can initiate and commit a change to Active Directory.
- **Pull replication.** A domain controller requests, or "pulls," changes from other domain controllers. As you learn more about replication, it may become easy to forget this, because a domain controller notifies its replication partners that it has changes to the directory, or a domain controller can poll its partners to see if they have changes to the directory. But the changes themselves are, in the end, requested or "pulled" by the target domain controller.
- **Store-and-forward replication.** A domain controller can pull changes from one partner, and then make those changes available to another partner. For example, domain controller B can pull changes initiated by domain controller A. Then, domain controller C can pull the changes from domain controller B.
- **Partitioning of the data store.** Domain controllers in a domain host only the domain naming context for their domain, which helps keep replication to a minimum, particularly in multidomain

forests. Other data, including application directory partitions and the partial attribute set (global catalog), are not replicated to every domain controller in the forest, by default.

- **Automatic generation of an efficient and robust replication topology.** By default, Active Directory will configure an effective, two-way replication topology so that the loss of any one domain controller does not impede replication. This topology is automatically updated as domain controllers are added, removed, or moved between sites.
- **Attribute-level replication.** When an attribute of an object is modified, only that attribute, and minimal metadata that describes that attribute, is replicated. The entire object is not replicated, except when the object is created.
- **Distinct control of intrasite replication** (within a single site) **and intersite replication** (between sites)**.** Replication can be distinctly controlled in both these situations.
- **Collision detection and management.** It is possible, although rare, that an attribute will have been modified on two different domain controllers during a single replication window. In such an event, the two changes will have to be reconciled. Active Directory has resolution algorithms that satisfy almost every such situation.

It is easier to understand Active Directory replication by examining each of its components. The following sections examine the components of Active Directory replication.

Intrasite Replication

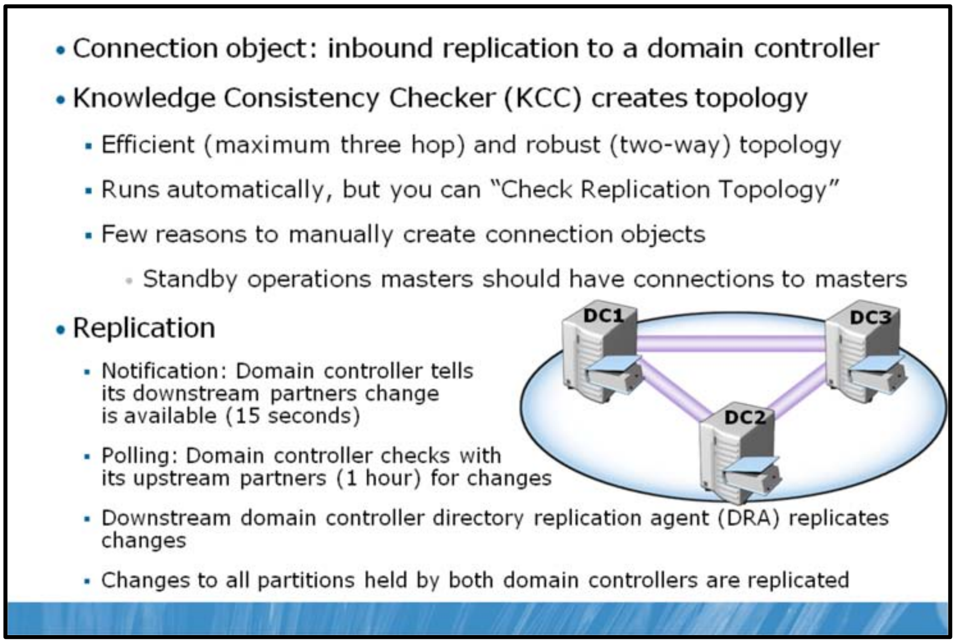

- Connection object: inbound replication to a domain controller
- Knowledge Consistency Checker (KCC) creates topology
  - Efficient (maximum three hop) and robust (two-way) topology
  - Runs automatically, but you can "Check Replication Topology"
  - Few reasons to manually create connection objects
    - Standby operations masters should have connections to masters
- Replication
  - Notification: Domain controller tells its downstream partners change is available (15 seconds)
  - Polling: Domain controller checks with its upstream partners (1 hour) for changes
  - Downstream domain controller directory replication agent (DRA) replicates changes
  - Changes to all partitions held by both domain controllers are replicated

There are several components of intrasite replication, which include:

- Connection Objects
- The Knowledge Consistency Checker
- Intrasite Replication
- Notification
- Polling

Connection Objects

A domain controller replicates changes from another domain controller because of AD DS connection objects, also called simply connection objects.

Connection objects appear in the administrative tools in the Active Directory Sites and Services snap-in as objects contained in the NTDS Settings container of a domain controller's server object.

The following screen shot shows an example: A connection object in SERVER02 configures replication from SERVER01 to SERVER02. A connection object represents a replication path from one domain controller to another.

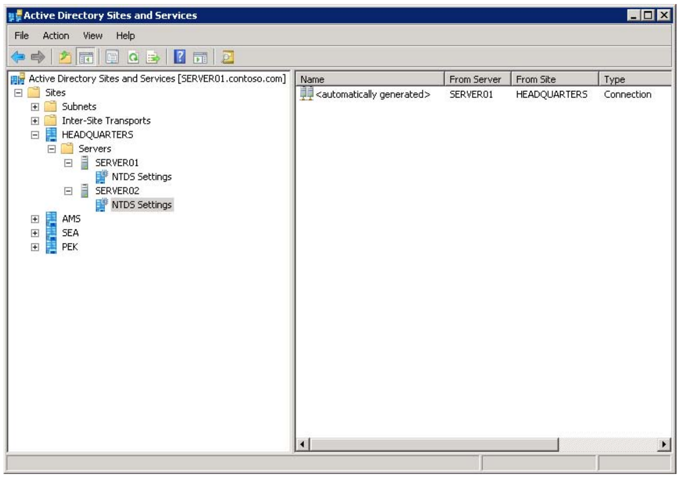

Name	From Server	From Site	Type
<automatically generated>	SERVER01	HEADQUARTERS	Connection

Connection objects are one-way, representing inbound-only replication. Replication in Active Directory is always a pull technology. In the domain illustrated here, SERVER02 pulls changes from SERVER01. SERVER02 is considered, in this example, a downstream replication partner of SERVER01. SERVER01 is the upstream partner. Changes from SERVER01 flow to SERVER02.

Note You can force replication between two domain controllers by right-clicking the connection object and selecting **Replicate Now**. Remember replication is inbound-only, so to replicate both domain controllers, you will need to replicate the inbound connection object of each domain controller.

The Knowledge Consistency Checker

The replication paths built between domain controllers by connection objects create the replication topology for the forest. Luckily, you do not have to create the replication topology manually. By default, Active Directory creates a topology that ensures effective replication. The topology is two-way so that if any one domain controller fails, replication continues uninterrupted. The topology also ensures that there are no more than three hops between any two domain controllers.

You'll notice in the previous screen shot that the connection object indicates it was automatically generated. On each domain controller, a component of Active Directory called the knowledge consistency checker (KCC) helps generate and optimize the replication automatically between domain controllers within a site. The KCC evaluates the domain controllers in a site and creates connection objects to build the two-way, three-hop topology described earlier. If a domain controller is added to or removed from the site, or if a domain controller is not responsive, the KCC rearranges the topology dynamically, adding and deleting connection objects to rebuild an effective replication topology.

You can manually create connection objects to specify replication paths that should persist. Manually created connection objects are not deleted by the KCC.

To create a connection object:

1. In **Active Directory Sites and Services**, locate the server object for the downstream replication partner—the domain controller that will receive changes from a source domain controller. Right-click the NTDS Settings container in the server object and click **New Active Directory Domain Services Connection**.

2. In the **Find Active Directory Domain Controllers** dialog box, select the upstream replication partner, and then click **OK**.
3. Give the new connection object a name, and then click **OK**.
4. Open the properties of the connection object; use the **Description** field to indicate the purpose of any manually created connection object.

Within a site, there are very few scenarios that would require creating a connection object. One such scenario is standby operations masters. Operations masters are discussed in Module 12.Select domain controllers as standby operations masters to be used in the event that the operations master role must be transferred or seized. A standby operations master should be a direct replication partner with the current operations master. Therefore, if a domain controller named DC01 is the RID master, and DC02 is the system that will take the RID master role if DC01 is taken offline, a connection object should be created in DC02 so that it replicates directly from DC01.

Intrasite Replication

After connection objects between the domain controllers in a site has been established—automatically by the KCC or manually—replication can take place. Intrasite replication involves the replication of changes within a single site.

Notification

Consider the site shown in the previous screen shot. When SERVER01 makes a change to a partition, it queues the change for replication to its partners. SERVER01 waits 15 seconds, by default, to notify its first replication partner, SERVER02, of the change. Notification is the process by which an upstream partner informs its downstream partners that a change is available. SERVER01 waits three seconds, by default, between notifications to additional partners. These delays, called the initial notification delay and the subsequent notification delay, are designed to stagger network traffic caused by intrasite replication.

Upon receiving the notification, the downstream partner, SERVER02, requests the changes from SERVER01, and the directory replication agent (DRA) performs the transfer of the attribute from SERVER01 to SERVER02. In this example, SERVER01 made the initial change to Active Directory. It is the originating domain controller and the change it made originates the change. When SERVER02 receives the change from SERVER01, it makes the change to its directory. The change is not called a replicated change, but it is a change nonetheless. SERVER02 queues the change for replication to its own downstream partners.

SERVER03 is a downstream replication partner of SERVER02. After 15 seconds, SERVER02 notifies SERVER03 that it has a change. SERVER03 makes the replicated change to its directory and then notifies its downstream partners. The change has made two hops, from SERVER01 to SERVER02 and from SERVER02 to SERVER03. The replication topology will ensure that there are no more than three hops before all domain controllers in the site have received the change. At approximately 15 seconds per hop, the change fully replicates in the site within one minute.

Polling

It is possible that SERVER01 might not make any changes to its replicas for quite a long time; particularly during off hours. In this case, SERVER02, its downstream replication partner, will not receive notifications from SERVER01. It is also possible that SERVER01 might be offline, which would also prevent it from sending notifications to SERVER02. So it’s important for SERVER02 to know that its upstream partner is online and simply does not have any changes.

This is achieved through a process called polling. Polling involves the downstream replication partner contacting the upstream replication partner with a query as to whether any changes are queued for replication. By default, the polling interval for intrasite replication is once per hour. It is possible, although not recommended, to configure the polling frequency from the properties of a connection object by clicking Change Schedule.

If an upstream partner fails to respond to repeated polling queries, the downstream partner launches the KCC to check the replication topology. If the upstream server is indeed offline, the site's replication topology is rebuilt to accommodate the change.

Site Links

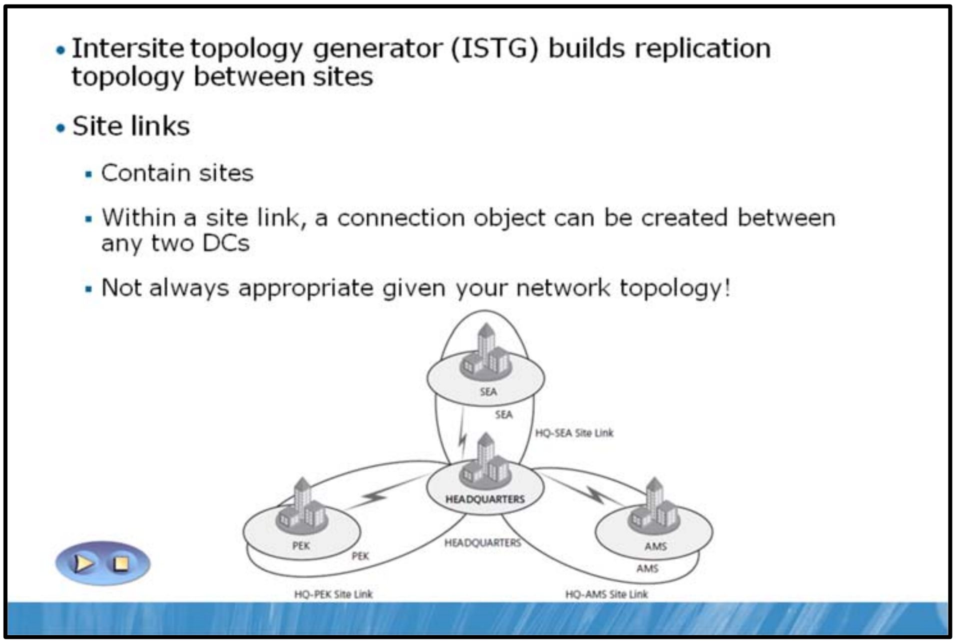

- Intersite topology generator (ISTG) builds replication topology between sites
- Site links
  - Contain sites
  - Within a site link, a connection object can be created between any two DCs
  - Not always appropriate given your network topology!

The KCC assumes that within a site, all domain controllers can reach each other. It builds an intrasite replication topology that is agnostic to the underlying network connectivity. Between sites, however, you can represent the network paths over which replication should occur by creating site link objects. A site link contains two or more sites. The intersite topology generator (ISTG), a component of the KCC, builds connection objects between servers in each of the sites to enable intersite replication—replication between sites.

Site links are greatly misunderstood, and the important thing to remember about a site link is that it represents an available path for replication. A single site link does not control the network routes that are used. When you create a site link and add sites to it, you are telling Active Directory that it can replicate between any of the sites associated with the site link. The ISTG creates connection objects, and those objects will determine the actual path of replication. Although the replication topology built by the ISTG effectively replicates Active Directory, it might not be efficient, given your network topology.

To better understand this concept, consider the following example. When you create a forest, one site link object is created: DEFAULTIPSITELINK. By default, each new site that you add is associated with the DEFAULTIPSITELINK. Consider an organization with a data center at the headquarters and three branch offices. The three branch offices are each connected to the data center with a dedicated link. You create sites for each branch office: Seattle (SEA), Amsterdam (AMS), and Beijing (PEK). The network and site topology is shown in the following figure.

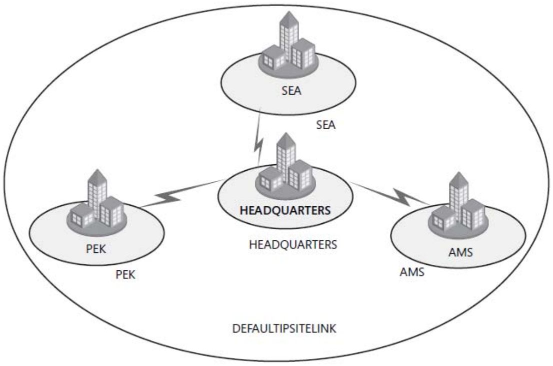

Because all four sites are on the same site link, you are instructing Active Directory that all four sites can replicate with each other. That means it is possible that Seattle will replicate changes from Amsterdam; Amsterdam will replicate changes from Beijing; and Beijing will replicate changes from the headquarters, which in turn replicates changes from Seattle. In several of these replication paths, the replication traffic on the network flows from one branch through the headquarters on its way to another branch. With a single site link, you do not create a hub-and-spoke replication topology even though your network topology is hub-and-spoke.

Therefore, manually create site links that reflect your physical network topology. Continuing the preceding example, you would create three site links:

- HQ-AMS, including the Headquarters and Amsterdam sites
- HQ-SEA, including the Headquarters and Seattle sites
- HQ-PEK, including the Headquarters and Beijing sites

Then, you would delete the DEFAULTIPSITELINK. The resulting topology is shown in the following figure.

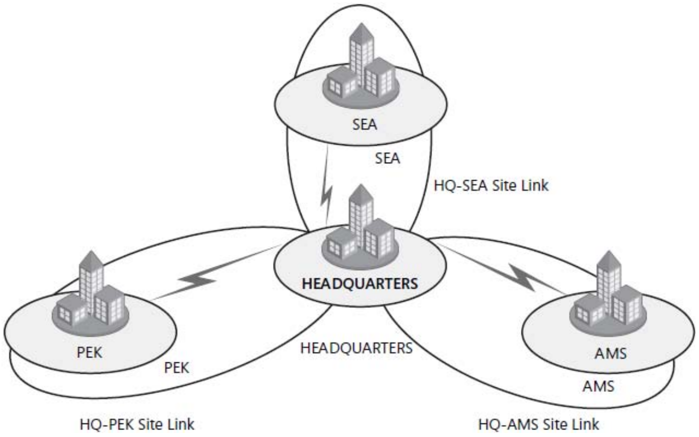

After you create site links, the ISTG will use the topology to build an intersite replication topology connecting each site. Connection objects will be built to configure the intersite replication paths. These connection objects are created automatically, and though you can create connection objects manually, there are few scenarios that require you to manually create intersite connection objects.

Replication Transport Protocols

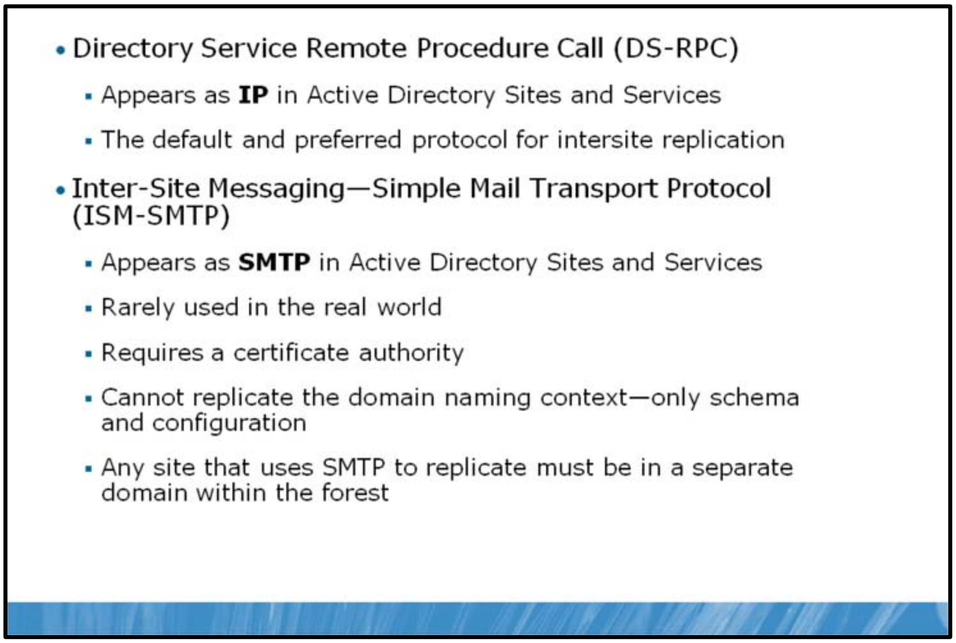

- Directory Service Remote Procedure Call (DS-RPC)
  - Appears as **IP** in Active Directory Sites and Services
  - The default and preferred protocol for intersite replication
- Inter-Site Messaging—Simple Mail Transport Protocol (ISM-SMTP)
  - Appears as **SMTP** in Active Directory Sites and Services
  - Rarely used in the real world
  - Requires a certificate authority
  - Cannot replicate the domain naming context—only schema and configuration
  - Any site that uses SMTP to replicate must be in a separate domain within the forest

In the Active Directory Sites and Services snap-in, you'll notice that site links are contained within a container named IP that itself is inside the Inter-Site Transports container. Changes are replicated between domain controllers by using one of two protocols:

Directory Service Remote Procedure Call (DS-RPC)

DS-RPC appears in the Active Directory Sites and Services snap-in as IP. IP is used for all intrasite replication and is the default, and preferred, protocol for intersite replication.

Inter-Site Messaging—Simple Mail Transport Protocol (ISM-SMTP)

Also known simply as SMTP, this protocol is used only when network connections between sites are unreliable or are not always available.

In general, you can assume you will use IP for all intersite replication. Very few organizations use SMTP for replication because of the administrative overhead required to configure and manage a certificate authority (CA), and because SMTP replication is not supported for the domain naming context. That is, a site uses SMTP to replicate to the rest of the enterprise, that site must be its own domain.

Bridgehead Servers

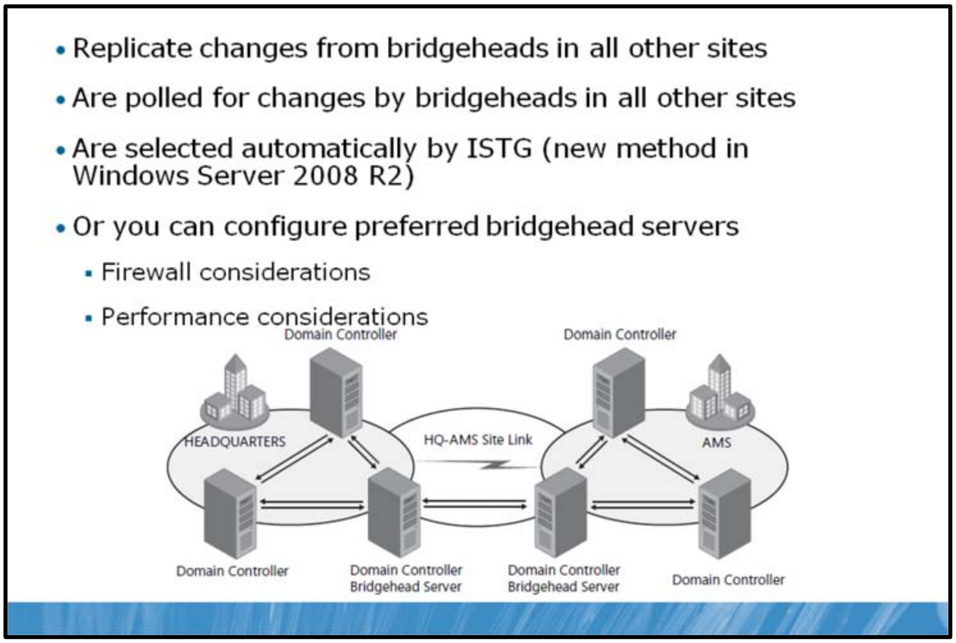

- Replicate changes from bridgeheads in all other sites
- Are polled for changes by bridgeheads in all other sites
- Are selected automatically by ISTG (new method in Windows Server 2008 R2)
- Or you can configure preferred bridgehead servers
  - Firewall considerations
  - Performance considerations

The ISTG creates a replication topology between sites on a site link. To make replication more efficient, one domain controller is selected to be the bridgehead server. The bridgehead server is responsible for all replication into and out of the site for a partition. For example, if a data center site contains five domain controllers, one of the domain controllers will be the bridgehead server for the domain naming context. All changes made to the domain partition within the data center will replicate to all domain controllers in the site. When the changes reach the bridgehead server, those changes will be replicated to bridgehead servers in branch offices, which in turn replicate the changes to domain controllers in their sites. Similarly, any changes to the domain naming context in branch offices will be replicated from the branches' bridgehead servers to the bridgehead server in the data center, which in turn replicates the changes to other domain controllers in the data center. The following figure illustrates intrasite replication within two sites and the intersite replication by using connection objects between the bridgehead servers in the sites.

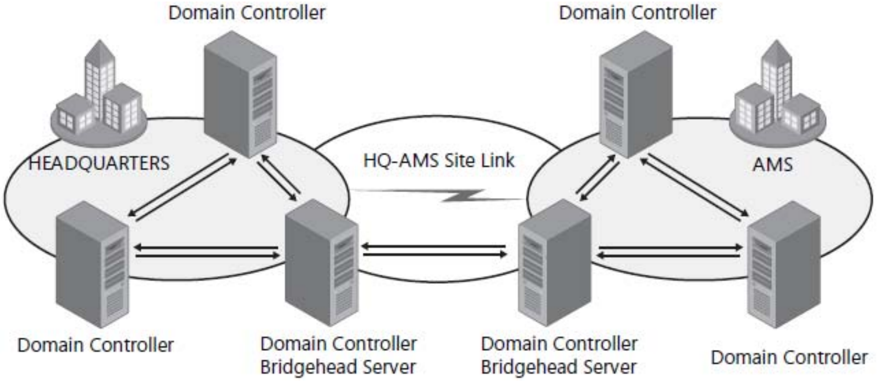

To summarize, the bridgehead server is the server responsible for replicating changes to a partition from other bridgehead servers in other sites. It is also polled by bridgehead servers in other sites to determine when it has changes that they should replicate.

Bridgehead servers are selected automatically, and the ISTG creates the intersite replication topology to ensure that changes are replicated effectively between bridgeheads sharing a site link. Bridgeheads are selected per partition, so it is possible that one domain controller in a site might be the bridgehead server for the schema and another might be for the configuration. However, you will usually find that one domain controller is the bridgehead server for all partitions in a site, unless there are domain controllers from other domains or application directory partitions, in which case, bridgeheads will be chosen for those partitions.

Preferred Bridgehead Servers

You can also designate one or more preferred bridgehead servers.

To designate a domain controller as a preferred bridgehead server:

1. Open the properties of the server object in the **Active Directory Sites and Services** snap-in.
2. Select the transport protocol, which will almost always be IP, and then click **Add**.

You can configure more than one preferred bridgehead server for a site, but only one will be selected and used as the bridgehead. If that bridgehead fails, one of the other preferred bridgehead servers will be used.

It's important to understand that if you have specified one or more bridgehead servers and none of the bridgeheads is available, no other server is automatically selected, and replication does not occur for the site even if there are servers that could act as bridgehead servers. In an ideal world, you should not configure preferred bridgehead servers. However, performance considerations might suggest that you assign the bridgehead server role to domain controllers with greater system resources. Firewall considerations might also require that you assign a single server to act as a bridgehead instead of allowing Active Directory to select and possibly reassign bridgehead servers over time.

Note The following content is specific to Windows Server 2008 R2.

In Windows Server 2008 R2, load balancing was introduced to distribute the workload evenly among bridgehead servers.

In pre–Windows Server 2008 R2 environments, inbound connections from sites flooded one domain controller in the hub site with requests. This was the case even if the connections to the hub site were in a load–balanced state. New bridgehead server selection technology improves the process of bridgehead server selection and avoids flooding a single server.

Site Link Transitivity and Bridges

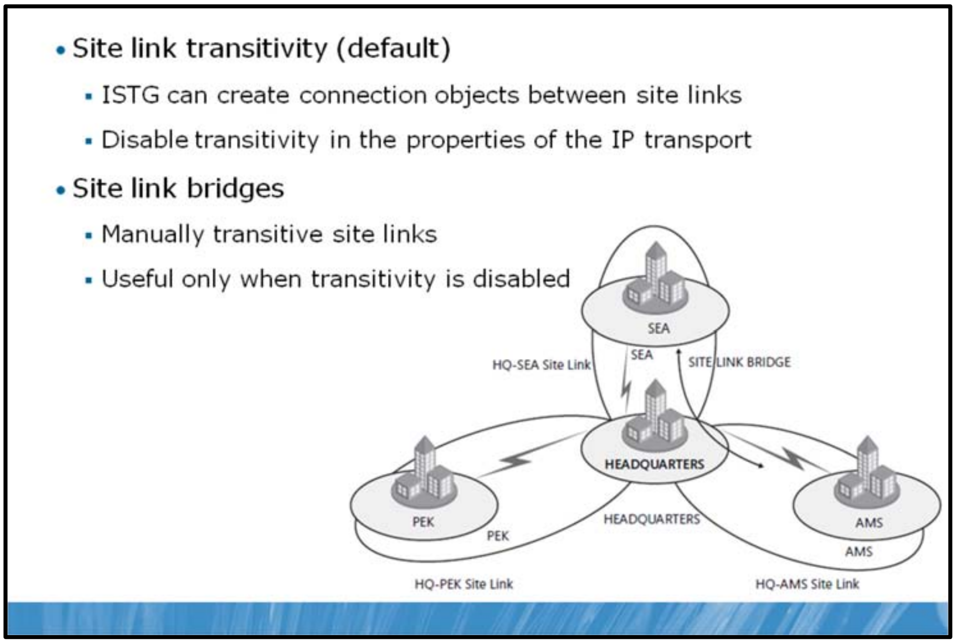

- Site link transitivity (default)
  - ISTG can create connection objects between site links
  - Disable transitivity in the properties of the IP transport
- Site link bridges
  - Manually transitive site links
  - Useful only when transitivity is disabled

After you have created site links and the ISTG has generated connection objects to replicate partitions between bridgehead servers that share a site link, your work might be complete. In many environments, particularly those with straightforward network topologies, site links might be sufficient to manage intersite replication. In more complex networks, however, you can configure additional components and replication properties.

Site Link Transitivity

By default, site links are transitive. If you consider the example from earlier, if the Amsterdam and Headquarters sites are linked, and the Headquarters and Seattle sites are linked, Amsterdam and Seattle will be transitively linked. This means, theoretically, that the ISTG could create a connection object directly between a bridgehead in Seattle and a bridgehead in Amsterdam, again working around the hub-and-spoke network topology.

You can disable site link transitivity by opening the properties of the IP transport in the Inter-Site Transports container and clearing the Bridge All Site Links check box. Before you do this in a production environment, read the technical resources about replication in the Windows Server technical libraries on Microsoft TechNet at http://go.microsoft.com/fwlink/?LinkID=214204.

Site Link Bridges

A site link bridge connects two or more site links in a way that creates a transitive link. Site link bridges are necessary only when you have cleared the Bridge All Site Links option for the transport protocol. Remember that site link transitivity is enabled by default, in which case, site link bridges have no effect.

The following figure illustrates the use of a site link bridge in a forest in which site link transitivity has been disabled. By creating a site link bridge, AMS-HQ-SEA, that includes the HQ-AMS and HQ-SEA site links, those two site links become transitive, so a replication connection can be made between a domain controller in Amsterdam and a domain controller in Seattle.

Control Intersite Replication

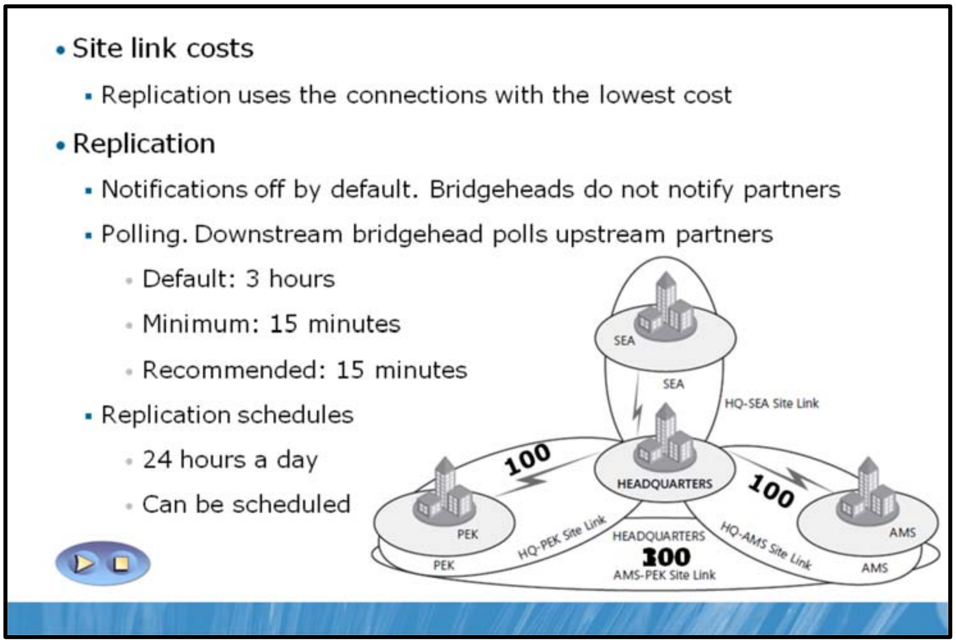

- Site link costs
  - Replication uses the connections with the lowest cost
- Replication
  - Notifications off by default. Bridgeheads do not notify partners
  - Polling. Downstream bridgehead polls upstream partners
    - Default: 3 hours
    - Minimum: 15 minutes
    - Recommended: 15 minutes
  - Replication schedules
    - 24 hours a day
    - Can be scheduled

Key points in this section include:

- Site Link Costs
- Replication Frequency
- Replication Schedules

Site Link Costs

Site link costs are used to manage the flow of replication traffic when there is more than one route for replication traffic. You can configure site link cost to indicate that a link is faster, more reliable, or is preferred. Higher costs are used for slow links, and lower costs are used for fast links. Active Directory replicates by using the connection with the lowest cost.

By default, all site links are configured with a cost of 100. To change the site link cost, open the properties of a site link and change the value in the **Cost** spin-box, shown in the following screen shot.

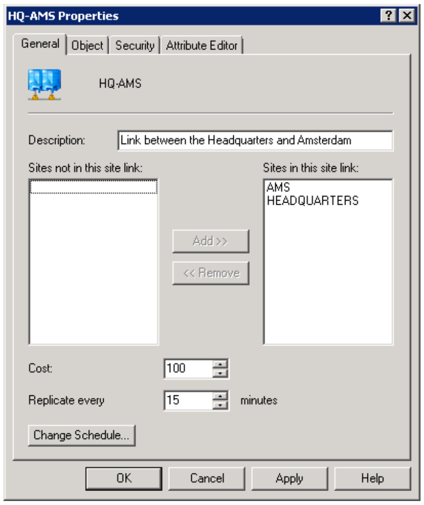

HQ-AMS

Description: Link between the Headquarters and Amsterdam

Sites not in this site link:

Sites in this site link:
AMS
HEADQUARTERS

Cost: 100

Replicate every 15 minutes

Returning to the example used earlier in the lesson, imagine if a site link was created between the Amsterdam and Beijing sites, as shown in the following figure.

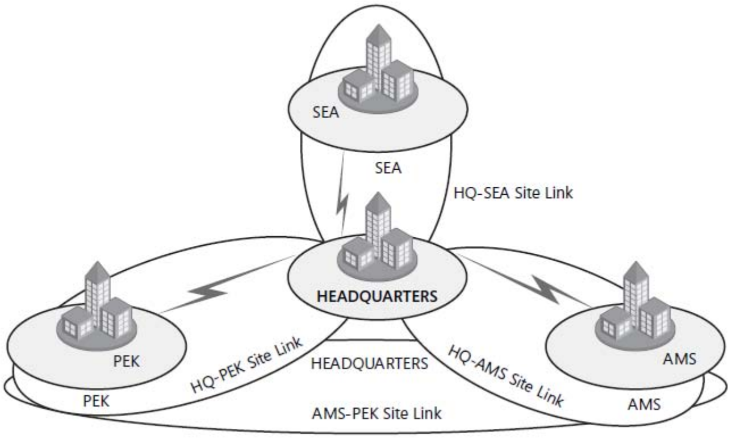

Such a site link could be configured to allow replication between domain controllers in those two sites if the links to the headquarters became unavailable. For example, you might want to configure such a topology as part of a disaster recovery plan.

With the default site link cost of 100 assigned to the AMS-PEK site link, Active Directory replicates changes directly between Amsterdam and Beijing. If you configure the site link cost to 300, changes replicate between Amsterdam and the Headquarters, then between the Headquarters and Beijing at a cost of 200 rather than directly over the AMS-PEK site link at a cost of 300.

Replication Frequency

Intersite replication is based only on polling; there is no notification. Every three hours, by default, a bridgehead server polls its upstream replication partners to determine whether changes are available. This replication interval is too long for organizations that want changes to the directory to replicate more quickly. You can change the polling interval for each site link.

To change the polling interval for a site link:

- Open the site link’s properties and change the value in the Replicate Every spin-box, shown in the previous screen shot.

The minimum polling interval is 15 minutes. That means, using Active Directory’s default replication configuration, a change made to the directory in one site will take several minutes before the change is replicated to domain controllers in another site.

Replication Schedules

By default, replication occurs 24 hours a day. However, you can restrict intersite replication to specific times by changing the schedule attributes of a site link. Open the properties of a site link and click Change Schedule. Using the Schedule For dialog box shown in the following screen shot, you can select the times during which the link is available for replication. The link shown in the figure does not replicate from 8:00 A.M. to 6:00 P.M. Monday through Friday.

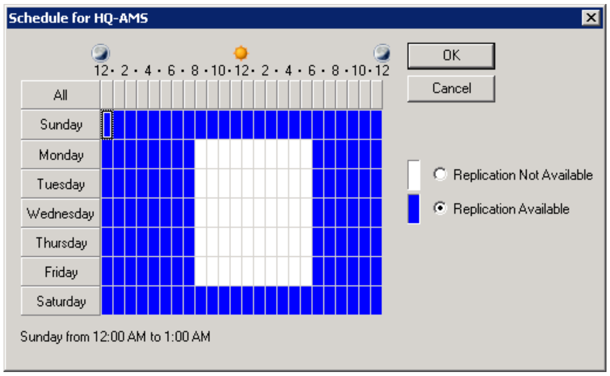

Schedule for HQ-AMS

○ Replication Not Available

● Replication Available

Sunday from 12:00 AM to 1:00 AM

You must be careful when scheduling site link availability. It is possible to schedule windows of availability that do not overlap, at which point, replication will not happen. Do not configure link availability. If you do not require link scheduling, select the Ignore Schedules option in the properties of the IP transport protocol. This option causes any schedules for site link availability to be ignored, ensuring replication 24 hours a day over all site links.

Monitor and Manage Replication

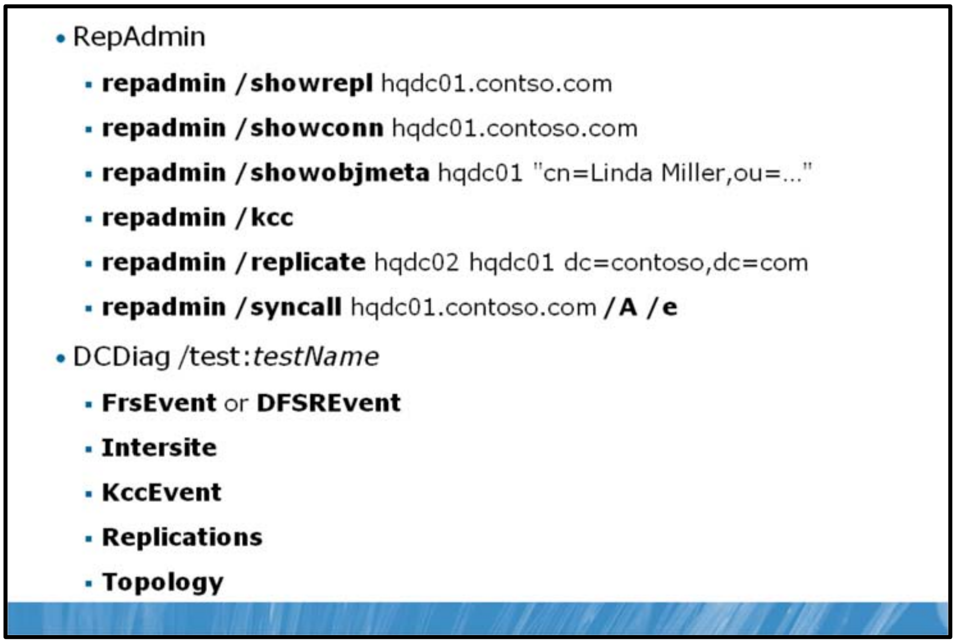

- RepAdmin
  - **repadmin /showrepl** hqdc01.contso.com
  - **repadmin /showconn** hqdc01.contoso.com
  - **repadmin /showobjmeta** hqdc01 "cn=Linda Miller,ou=..."
  - **repadmin /kcc**
  - **repadmin /replicate** hqdc02 hqdc01 dc=contoso,dc=com
  - **repadmin /syncall** hqdc01.contoso.com **/A /e**
- DCDiag /test:*testName*
  - **FrsEvent** or **DFSREvent**
  - **Intersite**
  - **KccEvent**
  - **Replications**
  - **Topology**

After you have implemented your replication configuration, you must be able to monitor replication for ongoing support, optimization, and troubleshooting. Two tools are particularly useful for reporting and analyzing replication: the Replication Diagnostics tool (Repadmin.exe) and Directory Server Diagnosis (Dcdiag.exe). This lesson introduces you to these tools.

Repadmin.exe

The Replication Diagnostics tool, Repadmin.exe, is a command-line tool that enables you to report the status of replication on each domain controller. The information produced by Repadmin.exe can help you spot a potential problem before it gets out of control and troubleshoot problems with replication in the forest. You can view levels of detail down to the replication metadata for specific objects and attributes, enabling you to identify where and when a problematic change was made to Active Directory. You can even use Repadmin.exe to create the replication topology and force replication between domain controllers.

Like other command-line tools, you can type repadmin /? to see the usage information for the tool. Its basic syntax is as follows.

```
repadmin command arguments...
```

Repadmin.exe supports a number of commands that perform specific tasks. You can learn about each command by typing repadmin /?:command. Most commands require arguments. Many commands take a *DSA_LIST* parameter, which is simply a network label (DNS, NetBIOS name, or IP address) of a domain controller. Some of the replication monitoring tasks you can perform by using Repadmin are:

- Display the replication partners for a domain controller. To display the replication connections of a domain controller, type repadmin /showrepl*DSA_LIST*. By default, Repadmin.exe shows only intersite connections. Add the /reps to argument to see intersite connections as well.

- Display connection objects for a domain controller. Type repadmin /showconn*DSA_LIST* to show the connection objects for a domain controller.
- Display metadata about an object, its attributes, and replication. You can learn a lot about replication by examining an object on two different domain controllers to find out which attributes have or have not replicated. Type repadmin /showobjmeta*DSA_LISTObject*, where *DSA_LIST* indicates the domain controller(s) to query. (You can use an asterisk [*] to indicate all domain controllers.) *Object* is a unique identifier for the object, its DN or GUID, for example.

You can also make changes to your replication infrastructure by using Repadmin. Some of the management tasks you can perform are:

- Launching the KCC. Type repadmin /kcc to force the KCC to recalculate the inbound replication topology for the server.
- Forcing replication between two partners. You can use Repadmin to force replication of a partition between a source and a target domain controller. Type repadmin /replicate*Destination_DSA_LIST Source_DSA_Name Naming_Context*.
- Synchronizing a domain controller with all replication partners. Type repadmin /syncall*DSA*/A /e to synchronize a domain controller with all its partners, including those in other sites.

Dcdiag.exe

The Directory Service Diagnosis tool, Dcdiag.exe, performs a number of tests and reports on the overall health of replication and security for Active Directory Domain Services. Run by itself, dcdiag.exe performs summary tests and reports the results. On the other extreme, dcdiag.exe /cperforms almost every test. The output of tests can be redirected to files of various types, including XML. Type dcdiag /? for full usage information.

You can also specify one or more tests to perform using the /test:*Test Name* parameter. Tests that are directly related to replication include:

- **FrsEvent.** Reports any operation errors in the file replication system (FRS).
- **DFSREvent.** Reports any operation errors in the DFS replication (DFS-R) system.
- **Intersite.** Checks for failures that would prevent or delay intersite replication.
- **KccEvent.** Identifies errors in the knowledge consistency checker.
- **Replications.** Checks for timely replication between domain controllers.
- **Topology.** Checks that the replication topology is fully connected for all DSAs.
- **VerifyReplicas.** Verifies that all application directory partitions are fully instantiated on all domain controllers hosting replicas.

See the Help & Support Center for more information about Repadmin.exe and Dcdiag.exe.

Lab B: Configure Replication

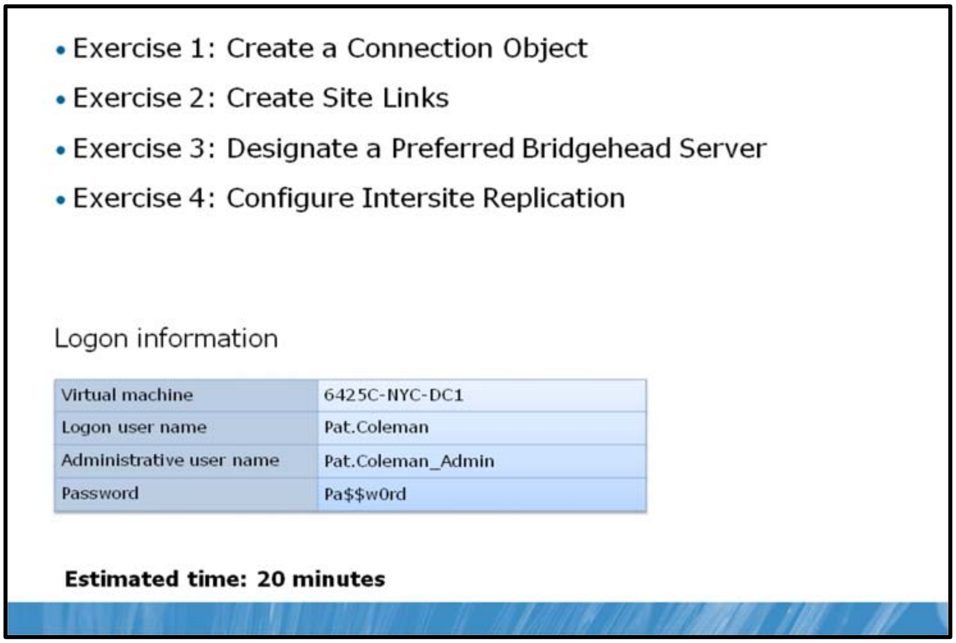

- Exercise 1: Create a Connection Object
- Exercise 2: Create Site Links
- Exercise 3: Designate a Preferred Bridgehead Server
- Exercise 4: Configure Intersite Replication

Logon information

Virtual machine	6425C-NYC-DC1
Logon user name	Pat.Coleman
Administrative user name	Pat.Coleman_Admin
Password	Pa$$w0rd

Estimated time: 20 minutes

Lab Setup

The virtual machines should already be started and available after completing Lab A. However, if they are not, complete Lab A first.

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V™ Manager, click **6425C-NYC-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:

- User name: **Pat.Coleman**
- Password: **Pa$$w0rd**
- Domain: **Contoso**

Lab Scenario

You are an administrator at Contoso, Ltd. You want to optimize replication of AD DS by aligning replication with your network topology and domain controller roles and placement.

Exercise 1: Create a Connection Object

It is a best practice to configure direct replication between a domain controller that will be a standby operations master and the domain controller that is currently the operations master. Then, if the current operations master needs to be taken offline, the standby operations master is as up to date as possible with the operations master. In this exercise, you will create a connection object between NYC-DC1 and NYC-DC2, where NYC-DC2, the standby operations master, replicates from NYC-DC1, the current operations master.

The main task for this exercise is as follows:

- Create a connection object.

▶ Task: Create a connection object.

1. Run **Active Directory Sites and Services** with administrative credentials. Use the account **Pat.Coleman_Admin** with the password **Pa$$w0rd**.
2. In the console tree, expand **HQ-BUILDING-2, Servers**, and **NYC-DC2**, and then click the **NTDS Settings** node below **NYC-DC2**.
3. Right-click **NTDS Settings** and click **New Active Directory Domain Services Connection**.
4. In the **Find Active Directory Domain Controllers** dialog box, select **NYC-DC1**, and then click **OK**, and answer **Yes** to the warning message.
5. In the **New Object – Connection** dialog box, type the name **NYC-DC1 - OPERATIONS MASTER**, and click **OK**.

   Question: Examine the properties of the connection object. Do not make any changes. What partitions are replicated from NYC-DC1? Is NYC-DC2 a GC server? How can you tell?

Results: In this exercise, you created a connection object to replicate changes from NYC-DC1 to NYC-DC2.

Exercise 2: Create Site Links

In this exercise, you will create site links between the headquarters and the other sites, creating a hub-and-spoke replication topology.

The main task for this exercise is as follows:

- Create site links.

▶ Task: Create site links.

1. Rename **DEFAULTIPSITELINK** to **HQ-HQB2**, and modify it so that it includes only the **HEADQUARTERS** and **HQ-BUILDING-2** sites.
2. Create a new IP site link named **HQ-BRANCHA** that includes the **HEADQUARTERS** and **BRANCHA** sites.

Results: In this exercise, you created two site links, one that links the HEADQUARTERS and HQ-BUILDING-2 sites, and one that links HEADQUARTERS and BRANCHA.

Exercise 3: Designate a Preferred Bridgehead Server

You can designate a preferred bridgehead server that will handle replication to and from its site. This is useful when you want to assign the role to a domain controller in a site with greater system resources or when firewall considerations require that the role be assigned to a single, fixed system. In this exercise, you will designate a preferred bridgehead server for the site.

The main task for this exercise is as follows:

- Designate a preferred bridgehead server.

▶ Task: Designate a preferred bridgehead server.

- Configure **NYC-DC2**as a preferred bridgehead server.

Results: In this exercise, you designated NYC-DC2as a preferred bridgehead server.

Exercise 4: Configure Intersite Replication

After you have created site links and, optionally, designated bridgehead servers, you can continue to refine and control replication by configuring properties of the site link. In this exercise, you will reduce the intersite replication polling frequency, and you will increase the cost of a site link.

The main task for this exercise is as follows:

- Configure intersite replication.

▶ Task: Configure intersite replication.

- Configure the replication interval for the **HQ-HQB2** site link to **15** minutes.
- Configure the replication interval for the **HQ-BRANCHA** site link to **15** minutes, and the cost to **200**.
- Examine the replication schedule for the **HQ-BRANCHA** site link. Experiment with configuring the schedule, but click **Cancel** after you are finished.

Results: In this exercise, you configured the intersite replication interval to 15 minutes for all site links.

▶ To prepare for the next module

When you finish the lab, revert the virtual machines to their initial state. To do this, complete the following steps:

1. On the host computer, start Hyper-V Manager.
2. Right-click **6425C-NYC-DC1** in the **Virtual Machines** list, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

Lab Review Questions

Question: Is the procedure you performed in Exercise 2 enough to create a "hub and spoke" replication topology, which ensures that all changes from branches are replicated to the headquarters before being replicated to other branches? If not, what should be done?

Module Review and Takeaways

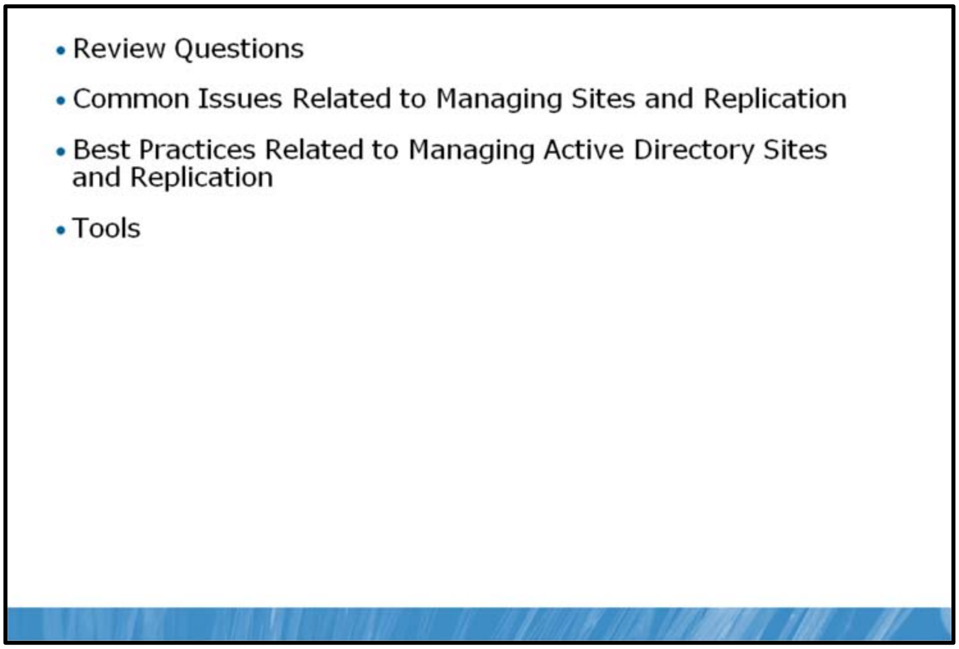

- Review Questions
- Common Issues Related to Managing Sites and Replication
- Best Practices Related to Managing Active Directory Sites and Replication
- Tools

Review Questions

Question: Why is it important that all subnets are identified and associated with a site in a multisite enterprise?

Question: What are the advantages and disadvantages of reducing the intersite replication interval?

Question: What is the purpose of bridgehead server?

Question: Which protocol can be used as an alternative for Active Directory replication? What is the disadvantage of using it?

Common Issues Related to Managing Sites and Replication

Issue	Troubleshooting tip
Client cannot locate domain controller in its site.	
Replication between sites does not work.	
Replication between two Domain Controllers in the same site does not work.	

Best Practices Related to Managing Active Directory Sites and Replication

You should implement the following best practices when you manage Active Directory sites and replication in your environment:

- Always provide at least one Global Catalog per site.

- Be sure that all sites have appropriate subnets associated.
- Do not setup long intervals without replication when you configure replication schedules for intersite replication.
- Avoid using SMTP as a protocol for replication.
- Do not use universal groups unless necessary because they create additional replication traffic.

Tools

Tool	Used for	Where to find it
Active Directory Sites and services	• Manage site objects • Manage site links • Manage replication	Administrative Tools
ADSI Edit	• View and manage Active Directory partitions	Administrative Tools
Repadmin	• Monitoring and managing replication	Command-line utility
dcdiag	Reports on the overall health of replication and security for Active Directory Domain Services	Command-line utility