



Operating System

Active Directory Site and Services Manager

Beta 3 Technical Walkthrough

Abstract

The primary purpose of the Microsoft® Active Directory™ Site and Services Manager is to administer the replication topology both within a site (in a local area network [LAN]) and between sites (in a wide area network [WAN]) in your enterprise environment.

Note: An Appendix following this technical walkthrough provides supporting background definitions and explanations on how Directory Service Replication is performed. If you are not familiar with replication, it may be appropriate for you to first review Appendix A.

© 1999 Microsoft Corporation. All rights reserved.

THIS IS PRELIMINARY DOCUMENTATION. The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This BETA document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Microsoft, Active Directory, the BackOffice logo, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other product or company names mentioned herein may be the trademarks of their respective owners.

*Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA
0499*

CONTENTS

INTRODUCTION	1
Sites	1
Default-First-Site-Name	1
USING THE SITES TOPOLOGY TOOL.....	4
Adding a Site	6
Adding a Subnet	7
SITE LINKS AND SITE LINK BRIDGES.....	9
Creating a Site Link	9
Creating a Site Link Bridge	10
BETA 3 IMPROVEMENTS.....	11
FOR MORE INFORMATION	12
Before You Call for Support	12
Reporting Problems	12
APPENDIX A– REPLICATION TOPOLOGY CONCEPTS	13
Sites	13
Servers and Windows 2000 Directory Service Settings	13
Connections	14
Site Links	14
Site Link Bridges	15

INTRODUCTION

The primary purpose of the Microsoft® Active Directory™ Site and Services Manager is to administer the replication topology both within a site (in a local area network [LAN]) and between sites (in a wide area network [WAN]) in your enterprise environment.

Note An Appendix following this technical walkthrough provides supporting background definitions and explanations on how Directory Service Replication is performed. If you are not familiar with replication, it may be appropriate for you to first review Appendix A.

Sites

A site is a region of your network with high bandwidth connectivity, and by definition is a collection of well connected machines—based on Internet Protocol (IP) subnets. Because sites control how replication occurs, changes made with this tool affect how efficiently domain controllers (DC) within a domain (but separated by great distances) will coalesce.

It is important to note that a site is a geographically determined boundary. This is separate in concept from Microsoft Windows® 2000 operating system domains, as a site may span multiple domains, and a domain may span multiple sites. Sites are not part of your domain namespace. Sites control replication of your domain information and help to determine resource proximity. For example, a workstation will select a DC within its site with which to authenticate.

To ensure that Active Directory, directory services in the Windows 2000 operating system, replicate properly, a service known as the Knowledge Consistency Checker (KCC) runs on all DCs and automatically establishes connections between individual machines in the same site. These are known as Windows 2000 Directory Service *connection objects*. An administrator may establish additional connection objects or remove connection objects, but at any point where replication within a site becomes impossible or has a single point of failure, the KCC will step in and establish as many new connection objects as necessary to resume Active Directory replication.

Replication between sites is assumed to occur on either higher cost or slower speed connections. As such, the mechanism for intersite (between site) replication permits the selection of alternative transports, and is established by creating Site Links and Site Link Bridges.

Default-First-Site-Name

Your first site was set up automatically when you installed Windows 2000 Server on the first domain controller in your enterprise. This is done by running DCPrmo.exe, and the resulting first site is called *Default-First-Site-Name*. This can be renamed later or left as is.

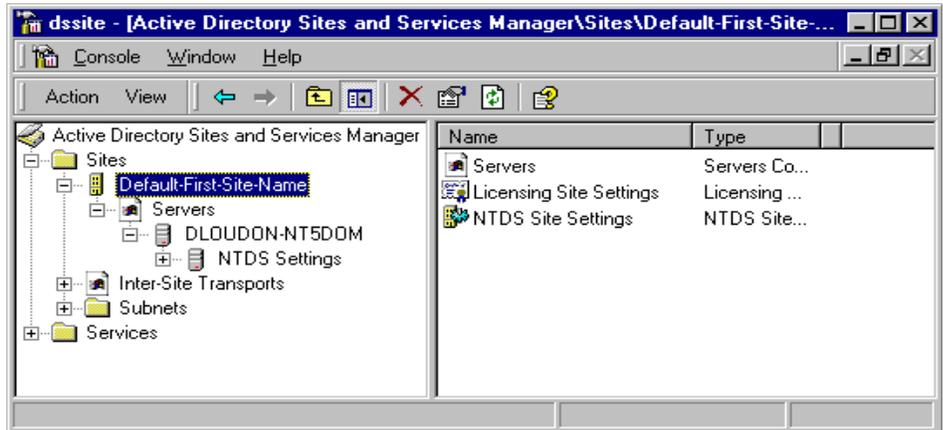


Figure 1: Final screen from DCPromo showing that a default first site name was created

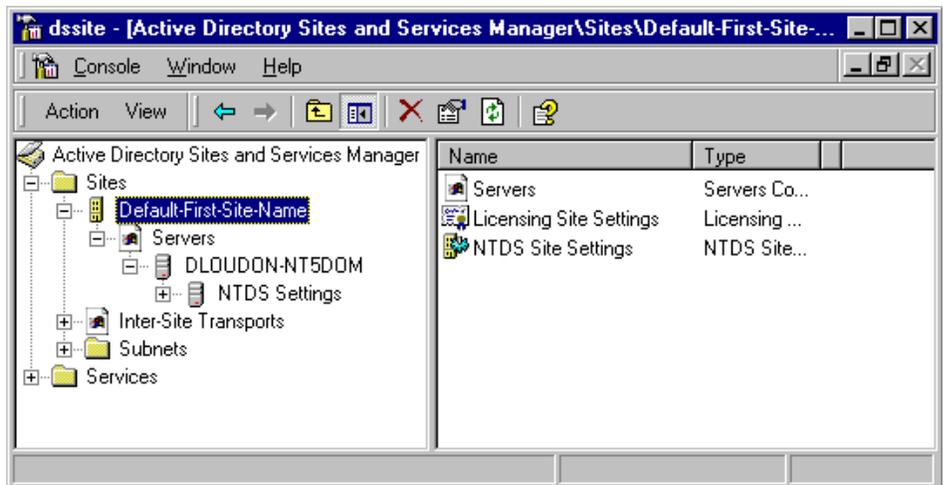


Figure 2: The resulting site information as displayed in Sites and Services Manager

The replication topology of sites on your network controls:

- **Where** replication occurs, such as which DCs communicate directly with which other DCs in the same site. Additionally, this topology controls how sites communicate with each other.
- **When** replication occurs. Replication between sites can be completely scheduled by the administrator. Replication between DCs inside the same site is notification based, where notifications are sent within five minutes of a change being made to an object in the domain.

All newly promoted Domain Controllers will be placed in the Site container that applies to them at time of install. For example, a server bound for California might have been initially built and configured in the Maui, Hawaii data center, therefore DC Promo will place the server in the Maui site. Once in California, the server object can be moved to the new site by way of the Sites and Services Manager.

You can use the sites portion of Sites and Services Manager to:

- Display the valid sites within an enterprise. As an example, *Default-First-Site-Name* might be a site name, as could be *Maui* or *Main-Campus*. You can create, delete, or rename sites.
- Display the servers which participate in a site. You can delete or move servers between sites. **Note:** Although you can also manually add servers, the task of adding a server is typically performed by DCPromo.)
- Display the applications that use site knowledge. The Active Directory Topology is rooted at *Sites\Default-First-Site-Name\Servers*. This contains just those servers participating in a specific site, regardless of domain. To view the connections for any given server, display *Sites\Default-First-Site-Name\Servers\{server}\NTDS Settings*. For each server, there are connections and schedules which control replication to other servers in this site.
 - *Connections*—For two machines to have two-way replication, a connection must exist from the first machine to the second, and a complimentary connection must exist from the second machine to the first.
 - *Schedules*—Within a site, pull replication of any new directory deltas occurs between servers approximately every five minutes. Schedules are significant within a site to force periodic notification to inbound partners in the event that a partner has a damaged connection object. This type of notification is typically every six hours. Additionally schedules are very significant in controlling pull replication between sites (there is no automatic five-minute replication between sites).
- Display transports and links between sites. *Transports* represent the protocols used to communicate between chosen sites (that is, IP).
- Display subnets. Subnets allow the administrator to associate ranges of IP addresses with sites.

USING THE SITES TOPOLOGY TOOL

To Run the Site and Services Manager Tool

1. From the Start menu, point to **Programs**.
2. Click **Administrative Tools**, and then click **Active Directory Sites and Services Manager**.

A console similar to the following appears (how similar depends on the nodes you've expanded and the names of your sites and servers).

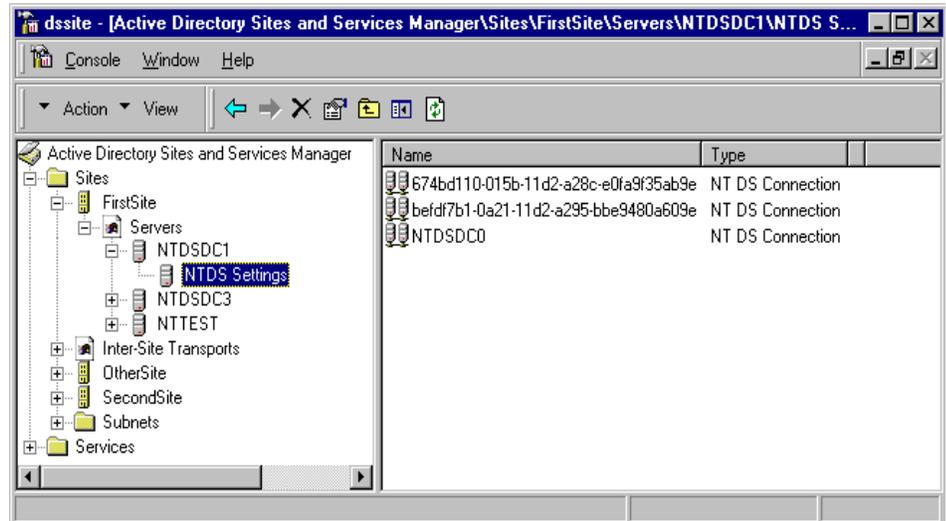


Figure 3: Sample console

The console above is focused under the computer NTDSDC1. Three connections are displayed. If these connections are established by the KCC, their names are globally unique identifiers (GUIDs).

The other computers likewise have two connections configured, so by using the above topology as a map, you could see the following:

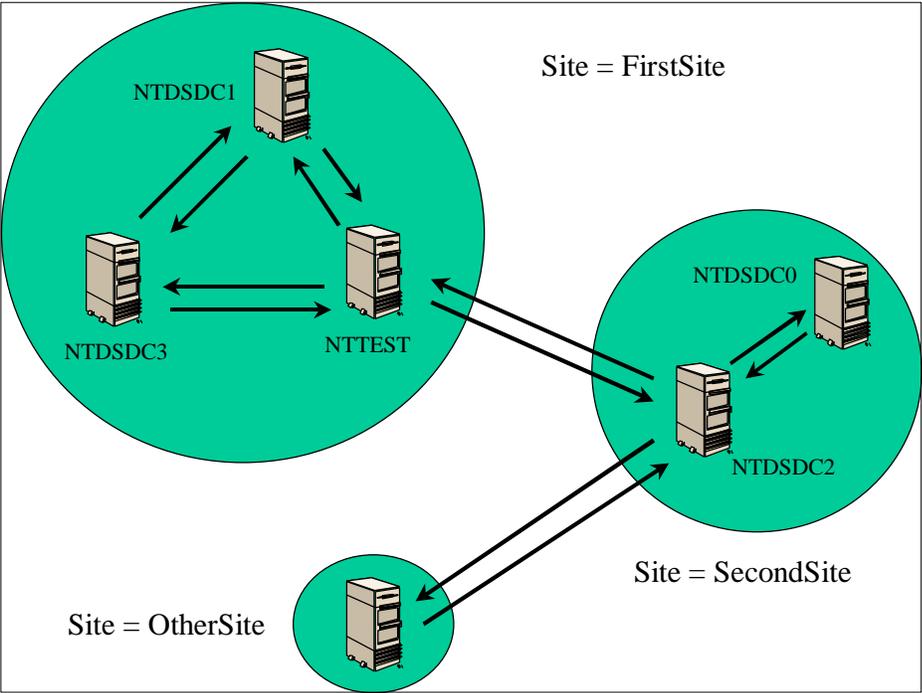


Figure 4. Sample topology

Adding a Site

To add new sites, use the Site and Services Manager

1. Right-click **Sites** in the tree in the left pane of the console, and then click **New**.
2. Click **Site**, and type in a name for the new site (*NewYork*, for example).
3. If presented with a Default Site Link, you might associate this site to a Site Link at this time. Site Links are explained later in this document.

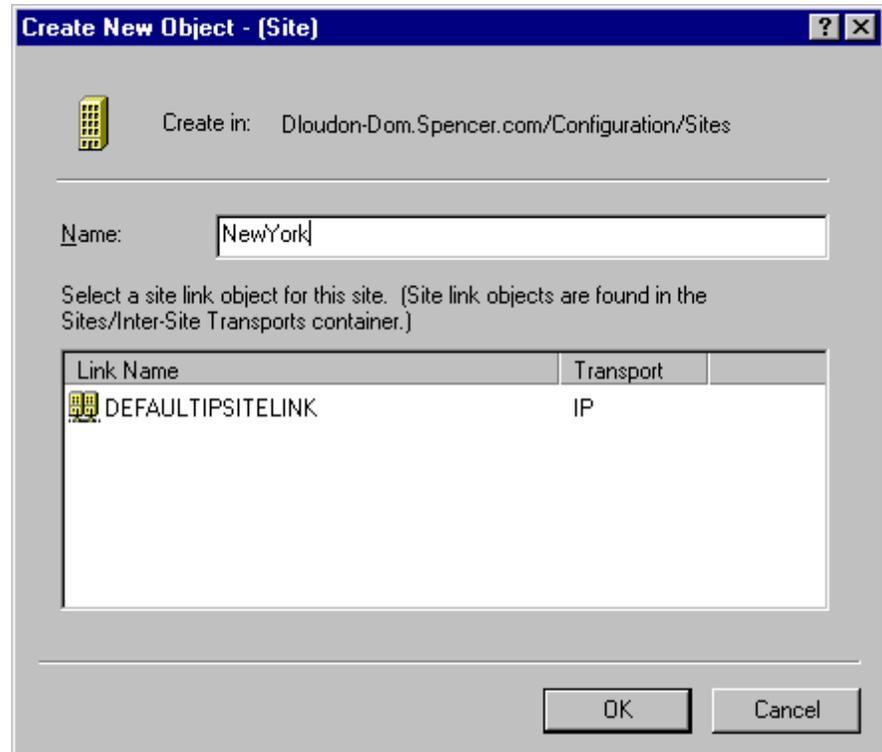


Figure 5. Creating a site

You can now move machines from other sites into this site (under the NTDS Settings container).

Adding a Subnet

To define subnets for a particular site

1. In the left pane of the console, right-click the **Subnets** item that appears under the site name.
2. From the **Action** menu, point to **New**, and then click **Subnets**.
3. In the **Name** box, type your subnet and subnet mask numbers (for example, 200.201.202.0/24).

If you have correctly entered the subnet, it will then appear in the Subnets folder.

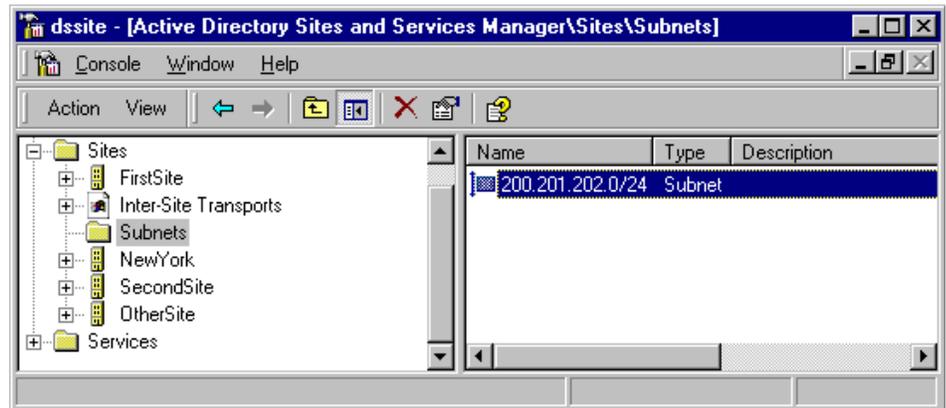


Figure 6. Adding a subnet

To associate the subnet with a site

1. Right-click the subnet in the right pane of the console, and then click **Properties**.
2. In the **Site** box, use the drop-down selection to select a site to associate with this subnet.

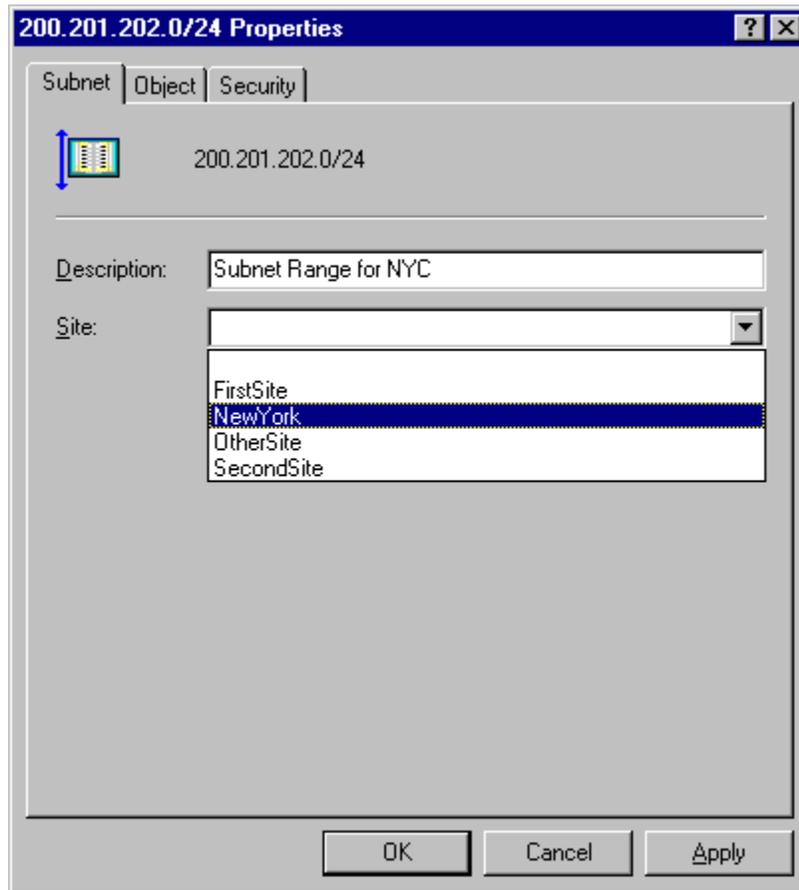


Figure 7. Associating a subnet with a site

SITE LINKS AND SITE LINK BRIDGES

Creating a Site Link

For scheduled replication to occur between multiple sites, both sites must agree on a transport to communicate. This will more than likely be IP based.

To create a link between two sites, perform the following:

1. From the Intersite Transports node, select one of the applicable transports. (Note that for Beta 3, Simple Mail Transfer Protocol (SMTP) or mailbased replication is not enabled.)
2. If you wish to join a site to an existing Site Link, select the link from the right hand panel, click **Properties** and add the Site.

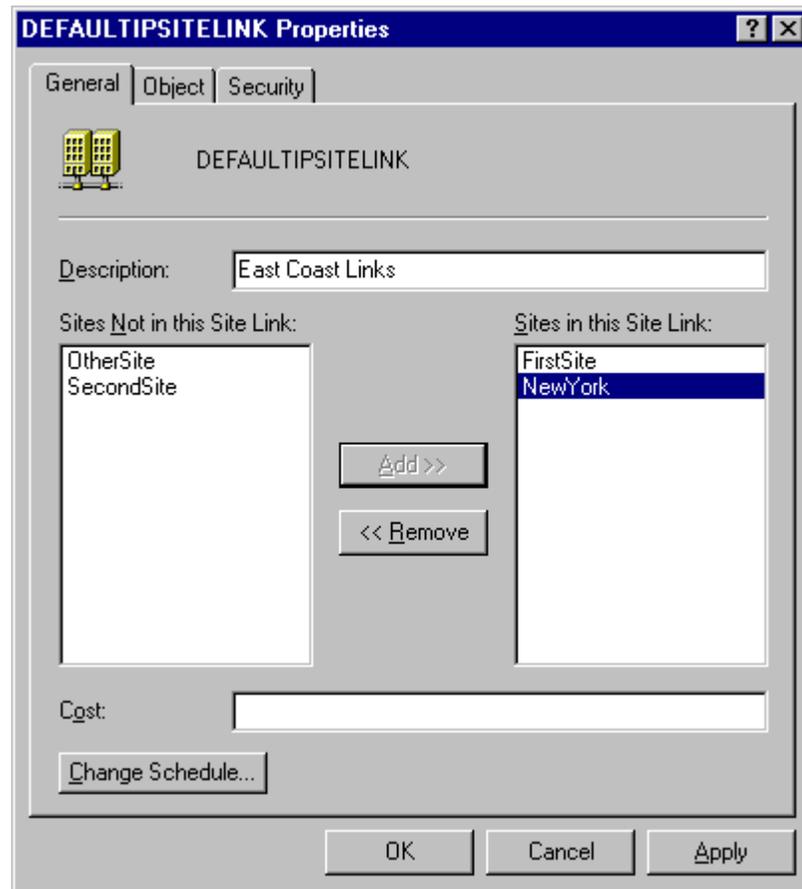


Figure 8. Joining a site to an existing link

3. Alternatively, to create a new Site Link, select **New, Site Link**, and complete the dialog box.

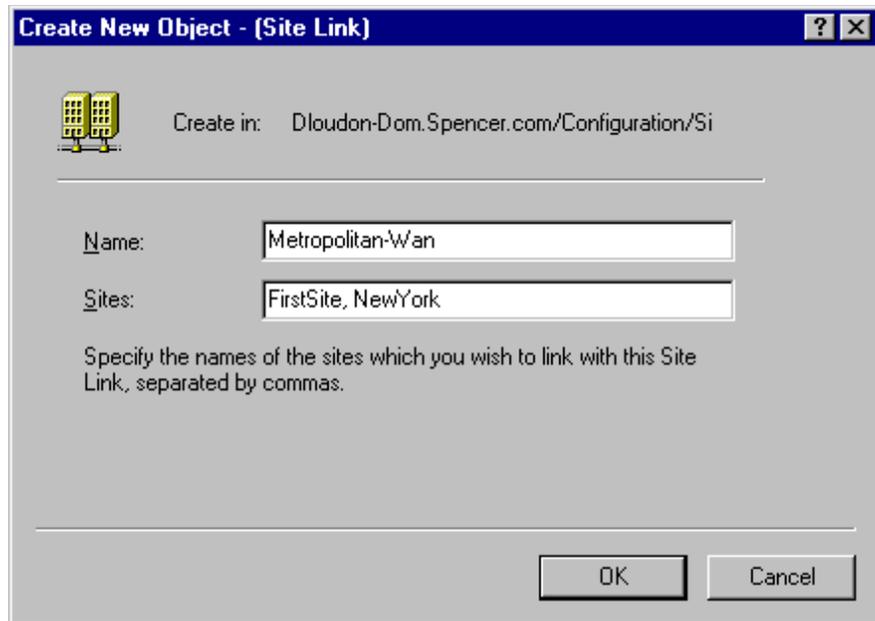


Figure 9. Creating a site link

Creating a Site Link Bridge

The process for creating a Site Link Bridge is identical to creating a Site Link; however, instead of providing Site names for the link, you're now providing Site link names for the bridge.

BETA 3 IMPROVEMENTS

Based on customer feedback, the following improvements are available in the Beta 3 release of Microsoft Windows® 2000:

- By default, all site links will be considered *transitive*. That is, all site links for a given transport will implicitly belong to a single site link bridge for that transport. So in the common case of a fully routed IP network, you won't have to configure *any* site link bridges. If your IP network is not fully routed you can turn off the transitive site link feature for the IP transport, in which case all IP site links will be considered intransitive and you'll configure site link bridges as in Beta 3 to model the actual routing behavior of your network.
- You'll be able to specify an additional value when creating a site link object: the *replication period*. If you don't specify this value, a global default replication period (that you can also set) will be assumed for the site link. When the KCC creates a connection object, its replication period will be the maximum of the periods along the minimum-cost path of site link objects from one end of the connection to the other.

As a result of this change, you can control topology and schedule independently:

- You control topology by setting the costs on the links. In a common scenario you might set cost = 1 for site links that are part of your backbone network, and cost = 100 for site links corresponding to slow connections to branch offices. Setting costs in this way ensures that a branch office replicates with a DC in a site that is part of the backbone, never directly with a second branch office. Unlike Beta 3, these cost numbers have no influence on the replication period.
- You control replication period by setting the replication period on site links. In a common scenario you might set the global default replication period to 15 minutes, and set a longer period on site links corresponding to slow connections to branch offices. The longer period makes more efficient use of the link but increases replication latency.
- You control link availability using the schedule on site links. You would use the default (100 percent available) schedule on most links, but might block replication traffic during peak business hours on links to certain branches. By blocking replication you give priority to other traffic but increase replication latency.
- You'll be able to enable connection objects for *change notification*. Change notification is automatic within a site and will remain so. But change notification will occur from a DC in one site to a DC in another if a connection object with change notification enabled, connects the two DCs in the correct orientation. The KCC will not create connections that are enabled for change notification, but an administrator can do so.

FOR MORE INFORMATION

For the latest information on Windows 2000, check out our World Wide Web site at <http://www.microsoft.com/ntserver> and the Windows NT Server Forum on the Microsoft Network (GO WORD: MSNTS).

For the latest information on the Windows 2000Beta 3, check out the World Wide Web site at <http://ntbeta.microsoft.com>

Before You Call for Support

Please keep in mind that Microsoft does not support these walkthroughs. The purpose of the walkthroughs is to facilitate your initial evaluation of Microsoft Windows 2000 features. For this reason, Microsoft cannot respond to questions you might have regarding specific steps and instructions.

Reporting Problems

Problems with Microsoft Windows 2000 Beta 3 should be reported by way of the appropriate bug reporting channel and alias. Please make sure to describe the problem adequately so that the testers and developers can reproduce it and fix it. Refer to the Release Notes included on the Windows 2000 Beta 3 distribution media for some of the known issues.

Sites

A *site* links a set of IP subnets. A *Site object* represents a site in Active Directory.

When you group a set of IP subnets into a site, you are saying that these subnets are well-connected. Intuitively, *well-connected* means high-bandwidth LAN connectivity (possibly involving hops through highperformance routers.)

Unfortunately, there is no cut-and-dried rule for deciding when two subnets should be placed in the same site. By understanding how Active Directory uses site information, you can make an informed decision. Active Directory exploits site information in the following four ways:

- When a client requests a connection to a Domain Controller (for example, for login), sites allow the client to connect to a DC within the same site whenever possible. This reduces network latency and conserves network bandwidth.
- When the Active Directory KCC configures replication connections between DCs, the KCC creates more connections between DCs in the same site than between DCs in different sites. The result is lower replication latency within a site, and lower replication bandwidth between sites.
- Replication messages between DCs in a site are uncompressed, so they use fewer CPU cycles on the DCs. Replication messages between DCs in different sites are compressed, so they use less network bandwidth.
- Replication between DCs in a site is triggered by the arrival of updates, reducing replication latency within a site. Replication between DCs in different sites is performed on a schedule, conserving network bandwidth.

Sites are not tied in any way to the Active Directory domain namespace. The name of a directory object does not reflect the site or sites in which the object is stored. A site may contain DCs from several domains, and DCs from a domain may be present in several sites. (In the Exchange Directory Service, sites are tied to the namespace.)

Servers and Windows 2000 Directory Service Settings

When you run DCpromo.exe to create a DC, the DC Promotion Wizard creates a *Server object* representing the DC's machine. A *Server object* is distinct from the *Computer object* that represents the machine as a security principal. In fact, the *Server object* contains a reference to the associated *Computer object*.

Server objects are children of *Site objects*. A server's parent site should contain the server's subnet. DCpromo cannot always place a *Server object* where it belongs; for example, the necessary *Site object* might not exist. In these cases, you must move a server from one site to another to keep the server's site consistent with its IP subnet.

A *NTDS Settings object* representing the Active Directory service running on the new DC is also created by DCpromo. The *NTDS Settings object* is the child of the DC's *Server object*.

Connections

A *Connection object* represents a replication connection from one DC to another. The Connection object is a child of the replication destination's NTDS Settings object, and points to the replication source.

Connection objects are created in two ways:

- By the KCC running on the destination DC.
- By a directory administrator.

A connection is unidirectional; a bidirectional replication connection is represented as two Connection objects under two different NTDS Settings objects.

Replication is performed between naming context (NC) replicas. Two DCs will often have several NCs in common. In fact, they *always* have at least two NCs in common: the Configuration NC and the Schema NC. If a connection exists from one DC to another, it will be used for replicating as many NCs as needed. There is never a need to create multiple connections linking the same two DCs in the same direction.

The KCC creates connections to keep your directory connected even in the case of extended failures and outages, with no manual intervention. Generally you create connections manually only if the KCC's automatically configured connections don't connect certain DCs that you believe should be connected:

- Within a site, you might decide to add connections in order to reduce the intra site replication latency. By default, an update takes at most three *hops* from where it originates in a site to any other DC in a site. To reduce this hop count to two or one, you would add extra links. The cost compared with the default configuration is extra CPU cycles, disk reads, and network messages spent on replication.
- Between sites, you might decide to add connections to reduce latency, especially in case of failures. (The KCC will create new connections to work around failures, but this adaptive process adds some latency). Here again the cost compared with the default configuration is extra CPU cycles, disk reads, and network messages spent on replication.

A connection includes a replication schedule. Another reason to create a connection manually is to achieve a replication schedule that can't be achieved by way of the KCC.

If a connection you create is identical to one the KCC would normally create, the KCC will not create an additional connection. And the KCC will never delete a connection you create.

Site Links

A *Site Link object* represents a set of sites that can communicate at uniform cost through some intersite transport. For IP transport, a typical site link connects just two sites and corresponds to an actual WAN link. An IP site link connecting more

than two sites might correspond to an asynchronous transfer mode (ATM) backbone connecting more than two clusters of buildings on a large campus, or several offices in a large metropolitan area connected through leased lines and IP routers.

You create a Site Link object for a specific intersite transport (typically IP transport) by specifying:

- A numeric cost. Higher cost numbers represent more expensive messages. And costs influence the frequency of replication on KCC-configured connections. Replication is every 15 minutes on a link with cost in [0..1], every 30 minutes on a link with cost in [1..2], every 45 minutes on a link with cost in [2..3], and so on. Cost units must be consistent for all site links in a directory.
- Two or more sites.
- A schedule. The schedule declares the time periods during which the link is available. For instance, you might make a site link for a dialup line unavailable during business hours when phone rates are high.

For example, if you create an IP site link object XYZ that connects three sites X, Y, and Z with cost 5, you are saying that an IP message can be sent between all pairs of sites (X to Y, X to Z, Y to X, Y to Z, Z to X, Z to Y) with cost 5.

If a site is common to two site links, that *does not* imply that the site will route between the two links. For instance, if site link XY connects sites X and Y through IP with cost 3, and site link YZ connects sites Y and Z through IP with cost 4, this does not imply that an IP message can be sent from X to Z with cost 7, or with any other cost for that matter.

A site can be connected to other sites by any number of Site Link objects. Each site in a multi-site directory must be connected by at least one site link. Otherwise it cannot replicate with DCs in any other site, so the directory is disconnected. Therefore, you must configure site links in a multisite directory.

Site Link Bridges

A *Site Link Bridge object* represents a set of site links, all of which can communicate by way of some transport. Typically a site link bridge corresponds to a route (or a set of routers) in an IP network.

You create a Site Link Bridge object for a specific intersite transport (typically IP transport) by specifying two or more site links for the specified intersite transport.

To understand what a site link bridge means, consider this example:

- Site link XY connects sites X and Y through IP with cost 3
- Site link YZ connects sites Y and Z through IP with cost 4
- Site link bridge XYZ connects XY and YZ.

The site link bridge XYZ implies that an IP message can be sent from site X to site Z with cost $3+4 = 7$. That is all the bridge does in this simple example.

Each site link L in a bridge should have some site in common with another site link in the bridge. Otherwise the bridge cannot compute the cost from sites in link L to the sites in other links of the bridge.

Multiple site link bridges for the same transport work together to model multihop routing. Add the following objects to the previous example:

- Site link WX connects sites W and X through IP with cost 2.
- Site link bridge WXY connects WX and XY.

Now the site link bridges WXY and XYZ together imply that an IP message can be sent from site W to site Z with cost $2+3+4 = 9$.

Any network that you can describe by a combination of site links and site link bridges, you can also describe by site links alone. But using site link bridges the network description is much smaller and easier for you to maintain, since you don't need a site link to describe every possible path between pairs of sites.