FOURTEEN

Understanding Active Directory Replication

In previous chapters, you have been introduced to Active Directory replication. Replication is the process of sending update information for data that has changed in the directory to other domain controllers. As a part of the Active Directory planning and implementation process, you should have a firm understanding of replication and how it takes place both within the domain and in multiple-site environments. This chapter provides you with a conceptual look at replication, both intrasite and intersite. You can learn more about intersite replication and sites in Chapter 15.

Active Directory Replication

Windows 2000 uses multi-master replication for the Active Directory. In multimaster environments, all domain controllers function as peers and all replicate Active Directory database changes to each other. There is no single master replicator, but all domain controllers are responsible for the replication tasks. Multi-master replication is effective because changes to the Active Directory can be made at any domain controller. The purpose of replication is to ensure that all domain controllers have accurate Active Directory data. For example, if an administrator adds a new object to the Active Directory on a particular domain controller, that domain controller is responsible for sending that change to all other domain controllers. Without effective replication, an Active Directory environment would quickly fall apart since each domain controller would be unaware of changes made by other domain controllers.

It is important to note the difference between directory replication and directory synchronization. Replication occurs between domain controllers in a Windows 2000 network. Directory synchronization occurs between different directories, such as Active Directory and Novell Directory Service (NDS). Since each directory uses a different schema, an agent is established that is considered a security principal for both directories. The agent replicates data between the two directory services by mapping between the two schemas. This process is known as synchronization.

The Replication Process

Active Directory replication is performed through multi-master replication and only changes are replicated. In other words, changes to the Active Directory can be made at any domain controller and only the change that is made will be replicated to all other domain controllers. The replication process is invisible to administrators and users.

Once a change has been made, the process ensures the data is replicated to domain controllers and that errors do not occur. The following sections outline the replication process.

Change Notification

The process begins with a "change notification." This change notification is sent to all domain controllers so they know there has been a change in the Active Directory database and that change is about to be replicated. Once the change notification has been sent, the process continues with an "update request."

Update Request

When a domain controller needs to replicate update data, an "originating update" is established. An originating update determines the kind of change that needs to be made to the Active Directory database. There are four different kinds of originating updates—add, modify, modifyDN, and delete. The add update adds an object to the Active Directory. For example, if you add a new shared folder, the add update replicates the information to all domain controllers. The modify update changes an attribute of an existing object. For example, if you change a user account's group membership, the modify update replicates this change. Next, the modifyDN update changes the name of an object or an object's parent. Finally, the delete update deletes an object from the Active Directory.

By using originating updates, changes to the Active Directory can be replicated (Figure 14.1). The task of creating an originating update is performed at the domain controller where the change was made. When the change is replicated to the other domain controllers, the update is known as a "replicated update" on those domain controllers. The replicated update occurs because of an originating update from another domain controller. When an originating update occurs, a stamp is attached to the updated attribute so it can be updated on all domain controllers.

A major feature of the replication process is Update Sequence Numbers (USNs). USNs are assigned numbers that are stored in a USN table on each domain controller. The USN table is used to determine the updates that need to occur between domain controllers. In other words, when a change occurs in the Active Directory, the domain controller where the change was made updates the USN so that all other domain controllers have an outdated USN for that attribute. When replication occurs, the USN is updated on all domain controllers. The USN allows other domain controllers to know they have an outdated USN and that the replication update

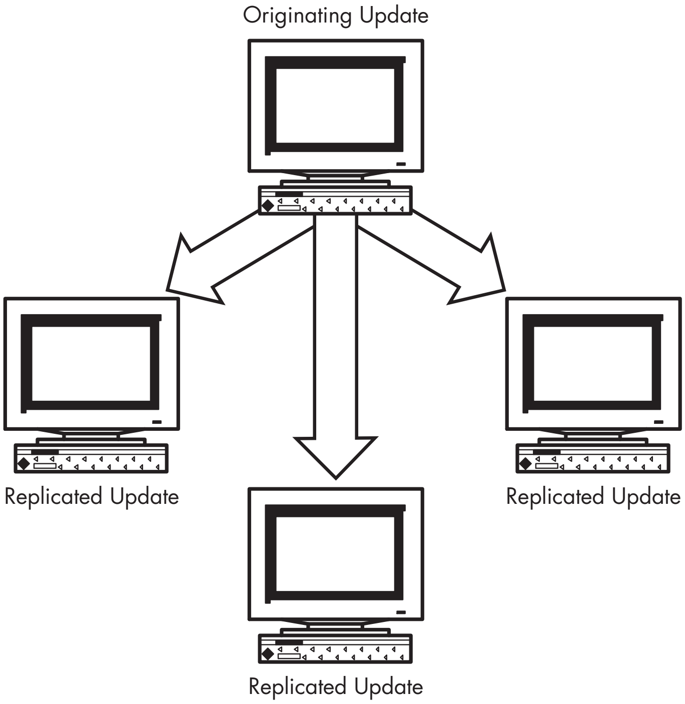

FIGURE 14.1 Originating and replicated updates

needs to be processed. Due to the use of USNs, timestamps on replication data are not necessary, although they are still maintained by the Active Directory and used in certain circumstances. For example, if an administrator at one domain controller makes a change to a user account phone number and an administrator at another domain controller makes the same change, the timestamp is used to "break the tie" between the two updates.

Another potential problem with Active Directory replication is unnecessary replication traffic. The Active Directory maintains a replication "loop" so that domain controllers have more than one path for sending and receiving replication traffic. However, the loop could allow updates to be sent to the same domain controller more than once. The Active Directory prevents this through a process called "propagation dampening." Propagation dampening allows domain controllers to detect when replication of data has already occurred on other domain controllers. When the domain controllers detect this, they do not send the same replication data where it has already been replicated. Figure 14.2 gives you a basic representation of propagation

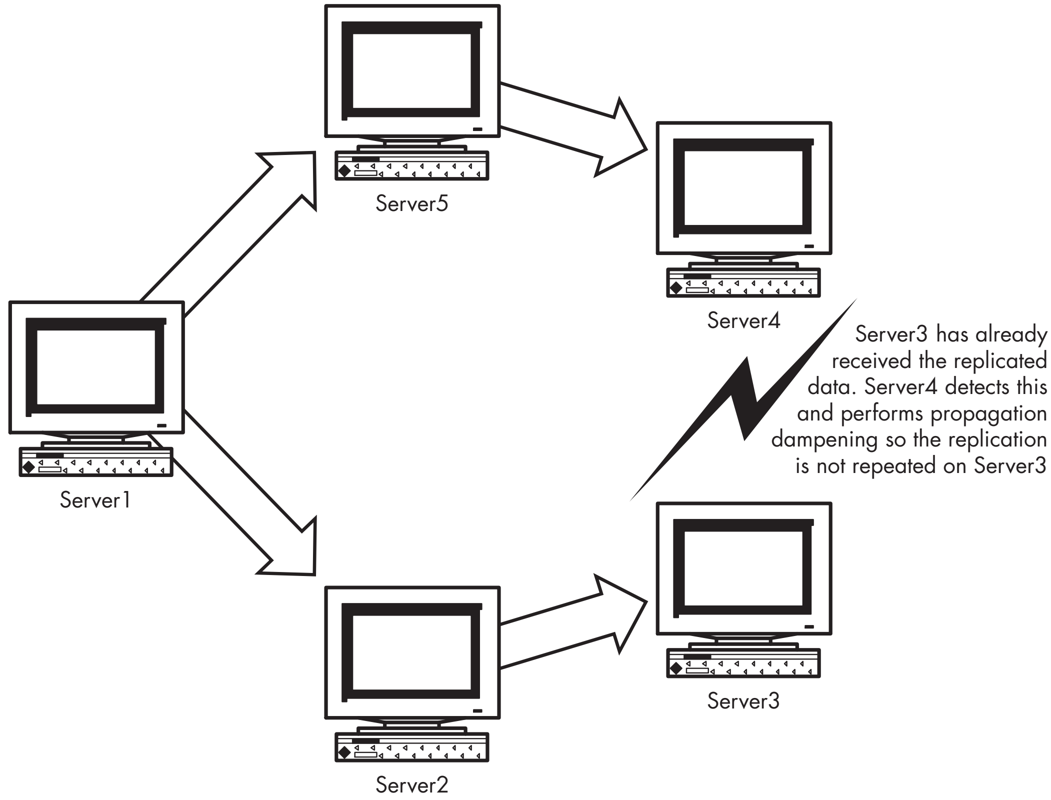

Server3 has already received the replicated data. Server4 detects this and performs propagation dampening so the replication is not repeated on Server3

FIGURE 14.2 Propagation dampening

dampening. Because replication has already occurred at Server3, Server4 halts the replication to Server3 through propagation dampening. This prevents Server3 from receiving the same replication data twice.

Propagation dampening occurs through the use of two vectors. Vectors are made up of pairs of data that combine a GUID (globally unique identifier) and a USN. The two vectors are called the up-to-date vector and the high watermark vector. The up-to-date vector contains server USN pairs and represents the highest originating update. The high watermark vector holds the USN numbers for attributes that have been added or modified in the directory and are stored in the replication metadata (which is simply "data about data") for that attribute. Through both vectors, propagation dampening can occur and unnecessary Active Directory updates can be avoided. Propagation dampening is an internal process and one that is invisible to administrators.

Along with propagation dampening, the Active Directory replication also has the task of solving replication conflicts. Since the Active Directory uses multi-master replication, there can be conflicts in changes. For example, if two administrators working on two different domain controllers make changes to the same attribute on the same object, a "collision" will occur when the data is replicated. The Active Directory attempts to minimize collisions by replicating data changes at the attribute level instead of the object level. This way, two different administrators on two domain controllers can make changes to the same object. As long as they are not making changes to the same attribute of that object, a collision will not occur. However, even the strategy of attribute level replication does not stop all collisions. In this case, the Active Directory must resolve collisions that occur. Resolution is accomplished through the use of timestamps and version numbers that are recorded in the metadata for that attribute. In the case of a collision, the domain controller(s) where the collision occurs will examine the timestamp and version number of the attribute and use the one that has the highest value. In the extreme case where the version numbers and timestamps match, then the highest Active Directory GUID would be used to break the tie. In any case, this collision resolution process ensures that the latest change to the object attribute always wins the resolution and is updated in the database.

Understanding Replication Partitions

Replication in an Active Directory environment functions at three major levels, and you should have a firm understanding of how partition replication occurs as this may impact some decisions you make as you design your Active Directory infrastructure. The three partitions are schema, configuration, and domain.

The schema partition contains object and attribute definitions. In other words, the schema partition contains a list of definitions that define what objects and attributes for those objects can exist in the Active Directory. Schema information is enterprise in nature—all domain controllers in a tree or forest share a common schema and any schema modifications are replicated across the forest. Because the schema defines objects and attributes, an object that is created, along with it's attributes, must conform to the definitions of the schema.

The configuration partition contains information about the physical structure of the Active Directory, such as the sites and domains and where domain controllers reside in the enterprise. Configuration information is replicated to all domain controllers in the tree or forest.

The domain partition contains information about all Active Directory objects that are specific to that domain, such as users and groups, OUs, and other resources. All domain partition information is completely replicated to all domain controllers within the domain. For global catalog servers in other domains, a read-only subset of the domain partition is replicated. This allows the global catalog server to know what is available in each domain so that other domain users can access resources, but changes to the domain partition can only be made from within the domain.

Obviously, a single Active Directory domain is much easier to implement in terms of replication. If your environment will use multiple domains, it is important to consider how replication will occur in your environment and how global catalog servers should be placed. You can learn more about this issue in Chapter 15.

Operations Masters

As stated several times, the Active Directory functions by using multi-master replication. Each domain controller can make changes to the Active Directory, and those changes are replicated. However, there are certain tasks that should only be performed on one domain controller due to the important nature of the tasks. In this case, the Active Directory makes certain only one domain controller can perform these tasks through a single-master approach. These single-masters, called operations masters, are used for the following tasks:

- Domain Operations Master—One domain controller in a forest can add or remove domains from the directory. This domain controller is called the domain operations master.
- PDC Emulator—One domain controller in the domain can function as the PDC emulator.

- Schema Operations Master—One domain controller in the forest can function as the schema operations master. Changes to the schema can only be performed on this domain controller, and schema modifications are replicated to all other domain controllers from the schema operations master.
- Infrastructure Operations Master—One domain controller in each domain can function as the infrastructure operations master. This domain controller can update SIDs of objects that move in and out of the domain.
- Relative Identifier (RID) Operations Master—One domain controller in the domain can be the RID operations master. The RID operations master can generate groups of Security Identifiers (SIDs) that are distributed to domain controllers in the domain.

Replication Topology

Now that you have a conceptual understanding of Active Directory replication, we turn our attention to the two types of replication topology—intrasite and intersite. Intrasite replication occurs within a site, which is a grouping of computers or domains that have high bandwidth. Typically, a site exists in a physical, geographic location, although sites can exist in multiple geographic locations. For example, a company has two sites, Dallas and Toronto. Each site contains one or more domains and exists in one geographic place where high bandwidth is available. In a site, the Active Directory can automatically generate a replication topology, or you can establish the replication topology if desired using the Active Directory Sites and Services tool (see Chapter 15). For intersite replication, replication occurs between two sites where there is normally lower bandwidth, or bandwidth is more expensive. You manually establish intersite replication through the Active Directory Sites and Services tool as it is not automatically generated by the Active Directory.

To fully understand replication topology, you need to understand how it is automatically accomplished in intersite communication through connection objects and the Knowledge Consistency Checker (KCC). The replication topology is the pathways domain controllers use to send and receive replication traffic.

Replication occurs through direct replication partners among domain controllers. The Active Directory determines whether a domain controller will function as a direct replication partner or whether it will receive replication data through transitive replication from other partners. This determination is made through the KCC. The KCC can determine how to establish the replication topology so that all domain controllers can receive replicated data. A connection object is defined as a potential direct replication partner (not transitive). Connection objects directly replicate with one another and transitive partners receive replication data indirectly. Connection objects are unidirectional and are established automatically by the Active Directory, or they can be established manually by an Active Directory administrator. The automatic topology generation, performed by the KCC, uses data you provide about sites and subnets within the site, as well as the cost of connections, to produce a replication topology in a bidirectional ring (by default). The ring is constructed so that the average number of hops a directory change will have to make is no more than three. When a change is made that needs to be replicated, the replication engine begins its job by waiting for an interval when replication can occur, which is every 5 minutes by default. When the interval arrives, the domain controller notifies its first replication partner. Then, all other partners are notified in a delay manner that you can configure, with the default being 30 seconds. So, in a typical site using the defaults, the propagation of the change to all domain controllers is 15 minutes or less.

As stated earlier, the Active Directory configures all of this on its own using information you enter in the Active Directory Sites and Services. You can manually configure the topology by creating additional connection objects or removing connection objects. The KCC does not delete any manually created connection objects. In the event of replication failure, however, the KCC will create new connection objects so that replication can resume.

In intrasite replication, the Active Directory normally does a good job of automatically generating a replication topology. For intersite replication, you will need to take a look at your existing links and the bandwidth available. You can also examine your network traffic to determine if changes in the replication topology need to be made. Network Monitor and Replication Monitor in Windows 2000 are good tools to use to examine network traffic.

Summary

This chapter provided you a conceptual look at replication in the Active Directory. Replication is performed using multi-master replication where each domain controller can replicate changes to the Active Directory to other domain controllers. Three replication partitions exist—schema, configuration, and domain—and an understanding of these partitions may affect decisions you make in your implementation. Replication topology functions on an intrasite and intersite level. The Active Directory automatically generates a replication topology for intrasite replication, and you can configure intersite replication using the Active Directory Sites and Services tool.

Case Study

As Bankston-Lewis continues to plan their Active Directory implementation, an issue of importance is training for network administrators. One of those training topics is replication. Administrators are all trained in Active Directory replication so that they all have an understanding of the processes and components that make up replication and replication topology. Since their environment consists of several sites and domains, replication configuration is of major importance. The administrators' general approach is to allow the Active Directory to automatically generate a replication topology for each site and make changes if necessary. Network Monitor and Replication Monitor can be used to examine replication traffic and determine if manual connection objects need to be established. For intersite communication, manual links will be established using the Active Directory Sites and Services tool.