CHAPTER 6

Building the Foundations for a Highly Available Architecture

Introduction

A lot of theory has been covered in the previous chapters. Now it's time to start implementing. This chapter and the ones to follow take what we have discussed up to now and roll it into an implementation plan for a data center that tens of thousands of users will rely upon.

First, we introduce clustering as it is accomplished on the Windows Server 2003 platform. We also discuss cluster concepts, models, and architecture. Then we implement the Active Directory architecture and network architecture as discussed in Chapter 5, "Preparing the Platform for a High-Performance Network," and lay the foundations for a highly available and reliable Web, database, and email server architecture, a network that will eventually comprise NLB IIS servers, NLB application servers, SQL Server clusters, Exchange clusters, and file and print clusters.

You can look at this chapter as the foundation implementation plan. It is what you need to follow if tasked with constructing and deploying a highly available solution. In the practical part, this chapter first outlines the process of building the forest and forest root domain, on either your lab or production network. It also covers the process of providing a resource for OS installations, tools, utilities, and patches. Then we prepare the cluster virtual server to begin hosting resources.

In this chapter, you implement Active Directory. At first glance, it seems that you are doing nothing more than setting up the usual AD network. But as you install the various cluster servers and services, you see that what is laid down in this chapter provides the solid foundation for the future systems. Then we deal with the actual process of clustering the servers, setting up cluster resources, and getting ready to activate the fail-over resources in Part II, "Building High Availability Windows Server 2003 Solutions." This is something you cannot do unless AD is well implemented beforehand.

Windows Clustering 101

There was a time in the not-too-distant past when the thought of clustering Windows servers sent a chill down the spines of network engineers and caused them to go take out long-term care insurance. Those days are gone with the clustering services that are now built into the Windows Server 2003 operating system. Only Windows Server 2003, Enterprise Edition and Windows Server 2003, Datacenter Edition can create clusters. Windows Storage Server 2003 is a version of Enterprise Edition for clustering file share resources.

There are two parts to clustering a high-availability service or application, such as Exchange 2003 or SQL Server 2000 or SQL Server 2005. The first part entails setting up the base cluster service and getting a virtual server going. The second part entails creating the resources that failover on that virtual server. Most of the second part of clustering is dealt with in Part II of this book.

By the time you are ready to cluster Exchange or SQL Server, you will be able to failover the virtual server resources from one node to the other and keep services, like drives and network interface cards, under the control of the cluster.

The Cluster Model

With Windows Server 2003, you have three models from which to choose; they are built into the operating system and are, thus, supported by Microsoft. The third option may require third-party software. Table 6.1 discusses the models in order of increasing complexity.

The most common cluster model (and inherited from Windows 2000 and Windows NT) is the single quorum cluster model in which multiple nodes of a cluster share a single quorum resource. In this model, all nodes communicate with each other across a local interconnect, and all nodes share a common disk array (in a SAN or a SCSI enclosure).

Windows Server 2003 also introduces the concept of a single node cluster, which is a cluster that is comprised of a single node or server. For obvious reasons, a single node cluster runs host cluster resources, but the cluster resources cannot fail-over to anything.

Then there is the geographic cluster or so-called "geo-cluster" in which the nodes that comprise the cluster are separated over a geographic divide. A wide area network usually separates the nodes and the geo-cluster nodes can be in different buildings or even across the country. They don't share storage or a quorum.

The central repository of data in a cluster is the so-called quorum resource. You can think of the quorum as the brain center of the cluster. The idea of a cluster is to provide system or server redundancy. In other words, when a server in the cluster fails, the cluster service is able to transfer operations to a healthy node. This is called failover. The quorum resource data is persistent and the quorum must survive node failure in the cluster or the resources cannot fail to the healthy node and start up.

This is why in a traditional, single quorum resource cluster, the quorum cannot be mounted into any single device on the node of the cluster unless the cluster can gain exclusive access to the device (and unless it can be moved or transferred upon node failure, which is technically possible even on a local disk resource as we will soon see). There are two exceptions to this rule: the single node cluster and the so-called geo-cluster, a concept in clustering now possible with Windows Server 2003.

Each of the cluster models discussed employs a different quorum resource type. Table 6.1 discusses the models.

Table 6.1 Cluster Model Options

Cluster Model	Application	Location of Cluster Configuration Data
Single Node	Ideal for labs, testing, development, and hosting applications on a virtual server	The quorum resource maintains the cluster configuration data either on a cluster storage device (an external drive array) or as a local drive on the node. Setup requires selection of the Local Quorum resource type.
Single Quorum	Typical local Active-Passive and Active-Active clusters	The quorum resource maintains the cluster configuration data on the single cluster storage device to which all nodes in the cluster are connected. Setup of this model ratifies the Physical Disk resource type (or other storage class resource type). The cluster installation will fail if this resource time does not test true as a viable quorum (we demostrate this later in the chapter).
Majority Node Set	Geographically dispersed server clusters	Geographic clusters are separated over wide area networks; therefore, each node maintains its own copy of the cluster configuration data. The quorum resource ensures the cluster configuration data is kept consistent across the nodes.

Single Node

Of particular interest is the Single Node cluster model in which the quorum resource can be maintained on a storage device on the local node. The idea behind the single node model is novel. With previous versions

of the operating system, it was impossible to establish a virtual server, what users attach to, on a cluster comprising only one node. The single node cluster enables this. The Single Node cluster model is illustrated in Figure 6.1.

NOTE: This chapter covers the creation of a single quorum cluster. However, we do touch on the subject of geo-clusters in Chapter 10, "High Availability, High-Performance Exchange."

You can use the single node cluster for lab testing of applications that have been engineered for clustering. You can also use it to test access to storage devices, quorum resources, and so on. The lab or development work is, thus, used to migrate the cluster-aware application into production as a standard single quorum cluster. It is also possible to simply cluster the single node with other nodes at a later time. The resource groups are in place and all you need to do is configure fail-over policies for the groups.

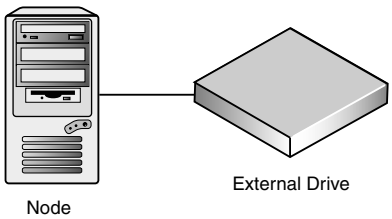

Figure 6.1 Single Node cluster model.

A single node cluster can also be used to simply provide a virtual server that users connect to. The virtual server service and name, thus, survives hardware failure. Both administrators and clients can see the virtual servers on the network and they do not have to browse a list of actual servers to find file shares.

What happens when the server hosting the single node cluster and the virtual server fail? The Cluster service automatically restarts the various application and dependent resources when the node is repaired. You can also use this service to automatically restart applications that would not otherwise be able to restart themselves.

For example, you can use this model to locate all the file and print resources in your organization on a single computer, establishing separate groups for each department. When clients from one department need to connect to the appropriate file or print share, they can find the share as easily as they would find an actual computer.

You can move the virtual server to a new node and end users never know the physical server behind the virtual server name has been changed. The real NetBIOS name of the server is never used. The downside of this idea is downtime. Moving the virtual server name to a new server requires downtime. Therefore, this is not suitable for a high-availability solution.

Single Quorum Cluster

This cluster model prescribes the quorum resource maintains all cluster configuration data on a single cluster storage device that all nodes have the potential to control. As mentioned earlier, this is the cluster model available in previous versions of Windows. The Single Quorum cluster model is illustrated in Figure 6.2.

Microsoft discounts the perception that the cluster storage device can be a single point of failure and promotes the idea that a *Storage Area Network* (SAN) where there are often multiple, redundant paths from the cluster nodes to the storage device mitigates in favor of this solution. While not discounting this model, if you study how a SAN is built, you discover there is some truth that a SAN is a single point of failure.

You can indeed have multple paths to the storage device (the "heart" of the SAN) as discussed in Chapters 3, "Storage for Highly Available Systems," and 4, "Highly Available Networks." However, the SAN controller is really nothing more than a server with an operating system that is dedicated to hosting the drive arrays in its enclosures. Unless you have redundant controllers, your SAN will fail if a component in the SAN controller fails. SAN memory can fail, its operating system can hang, the processors can be fried, and so on. Thus, to really elimimate every single point of failure in this model, you really need to have two SANs on the back end. This idea really opens a can of worms. After all, most IT shops do not budget for two SANs for every cluster. The SANs of today have many redundant components within their single footprint (usually a very large footprint) in the data center. To deploy two-mirrored SANs on a cluster is not only a very expensive proposition, but it is technically very difficult to install and manage.

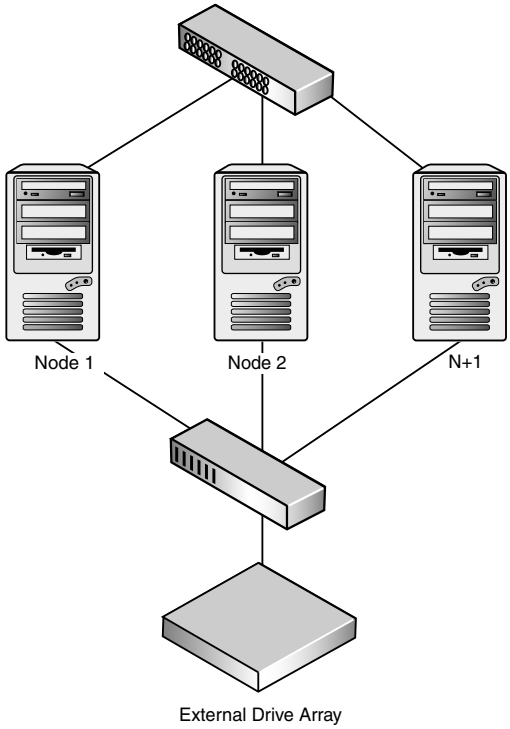

Figure 6.2 Single Quorum cluster model.

Majority Node Set

As mentioned, geo-cluster nodes can reside on opposite sides of the planet because each node maintains its own copy of the cluster configuration data. The quorum resource in the geo-cluster is called the Majority Node Set resource. Its job is to ensure the cluster configuration data is kept consistent across the different nodes; it is essentially a mirroring mechanism. The Majority Node Set cluster model is illustrated in Figure 6.3.

The quorum data is transmitted unencrypted over *Server Message Block* (SMB) file shares from one node to the other. Naturally, the cluster nodes cannot be connected to a common cluster disk array, which is the main idea behind this model.

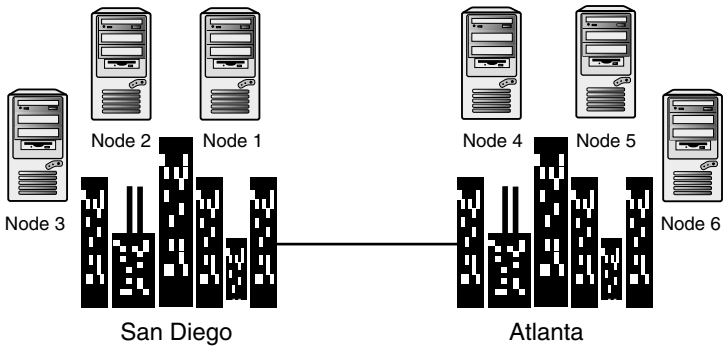

Figure 6.3 Majority Node Set cluster model.

You can use a majority node set cluster in special situations, and it will likely require special third-party software and hardware offered by your *Original Equipment Manufacturer* (OEM), *Independent Software Vendor* (ISV), or *Independent Hardware Vendor* (IHV).

Let's look at an example. Let's say we create an 8-node geo-cluster. We could, for example, locate four nodes in one data center, say in Atlanta, and the other 4 nodes in another data center in Phoenix. This can be achieved, and you can still present a single point of access to your clients. At any time a node in the geo-cluster can be taken offline, either intentionally or as a result of failure, and the cluster still remains available.

You can create these clusters without cluster disks. In other words, you can host applications that can failover, but the data the application needs are replicated or mirrored to the quorum data repositories on the other nodes on the cluster. For example, we can use this model with SQL Server to keep a database state up-to-date with log shipping. In Chapter 10, we investigate the particular solutions offered by NSI Software: Double-Take and GeoCluster.

The majority node set is enticing, but there are disadvantages. For starters, if more than half the nodes fail at any one time, then the entire cluster itself fails. When this happens, we say the cluster has lost quorum. This fail-over limitation is in contrast to the Single Quorum cluster model discussed ealier which will not fail until the last node in the cluster fails.

The Quorum Resource

Every cluster requires a resource which is designated as the quorum resource. The idea of the quorum is to provide a place to store configuration data for the cluster. Thus, when a cluster node fails, the quorum lives to service the new active node (or nodes) in the cluster. The quorum essentially maintains the configuration data the cluster needs to recover.

This data in the quorum is saved in the form of recovery logs. These logs store the changes that have been saved in the cluster database. Each node in the cluster depends on the data in the cluster database for configuration and state.

A cluster cannot exist without the cluster database. For example, a cluster is created when each node that joins the cluster updates its private copy of the cluster database. When you add a node to the existing cluster, the Cluster service retrieves data from the other active nodes and uses it to expand the cluster. When you create the first node in a cluster, the creation process updates the cluster database with details about the new node. This is discussed in more detail in the section "Clustering" later in this chapter.

The quorum resource is also used by the cluster service to ensure the cluster is composed of an active collection of communicating nodes. If the nodes in the cluster can communicate normally with each other (across the cluster interconnect), then you have a cluster. Like all service databases on the Windows platform, the cluster database and the quorum resource logs can become corrupt. There are procedures to fix these resources and we cover this a little later in this chapter.

When you attempt to create a cluster, the first node in the cluster needs to gain control of the quorum resource. If it cannot see the resource (this quorum), then the cluster installation fails. We show you this later. In addition, a new node is allowed to join a cluster or remain in the cluster only if it can communicate with the node that controls the quorum resource.

Let's now look at how the quorum resource is used in a two-node cluster, which is the type of cluster we will in the coming chapters.

When the first node in the cluster fails, the second node continues to write changes to the cluster database that it has taken control of. When the first node recovers and a fail-back is initiated, then ownership on the cluster database and quorum resource is returned in the fail-back mechanism.

But what if the second node fails before the first is recovered? In such a case, the first node must first update its copy of the cluster database with the changes made by the second node before it failed. It does this using the quorum resource recovery logs.

If the event the interconnect between the nodes fails, then each node automatically assumes the other node has failed. Typically, both nodes then attempt to continue operating as the cluster, and what you now have is a state called *split brain sydrome*. Imagine both servers succeeded in operating the cluster, you would then have two separate clusters claiming the same virtual server name and competing for the same disk resources. This is not a good condition for a system to find itself in.

The operating system prevents this scenario with quorum resource ownership. The node that succeeds in gaining control of the quorum resource wins and continues to present the cluster. In other words, whoever controls the "brain" wins. The other node submits, the fail-over completes, and the resources on the failed node are deactivated.

What constitutes a valid quorum resource? The quorum can be any resource that meets the following attributes:

- It can be accessed by a single node that must be able to gain physical control of it and defend the control.
- It must reside on physical storage that can be accessed by any node in the cluster.
- It must be established on the NTFS file system.

It is possible to create custom resource types as long as developers meet the arbitration and storage requirements specified in the API exposed by the Microsoft Software Development Kit. Let's now look at some deployment scenarios.

Deployment Scenarios

Let's discuss some example deployment schemes, namely the *n*-node fail-over scheme, the fail-over ring scheme, and the hot-standby server scheme.

In the *n*-node fail-over scheme you deploy applications that are setup to be moved to a passive node when the primary node on a 2-node cluster fails. In this configuration, you limit the possible owners for each resource group. You will see how we do this in Part II of this book.

Let's consider the so-called N+I hot-standby server scheme. Here you reduce the overhead of the 2-node failover by adding a "spare" node (one for each cluster pair) to the cluster. This provides a so-called "hot-standby" server that is part and parcel of the cluster and equally capable of running the applications from each node pair in the event of a failure of both of the other nodes. Both of these solutions are called active/passive clusters—*n*-node and *n*-node+1 (or N+1).

As you create the N+1 mode cluster, you will discover it is a simple matter to configure as the spare node. How you use a combination of the preferred owners lists and the possible owners list depends on your application. You typically set the preferred node to the node that the application runs on by default; and you set the possible owners for a given resource group to the preferred node and the spare node.

Then there is the concept of a *Failover Ring*. Here you set up each node in the cluster to run an application instance. Let's assume we have an instance of SQL Server on each node of the cluster. In the event of a failure, the SQL Server on the failed node is moved to the next node in sequence. Actually, an instance of SQL Server is installed on every server. Fail-over simply activates the SQL Server instance, and it takes control of the databases stored on the SAN or SCSI array. We call this the Active-Active cluster.

You can also allow the server cluster to choose the failover node at random. You can do this with large clusters and you'll just not define a preferred owners list for the resource groups. In other words, each resource group that has an empty preferred owners list is failed over to any node in random fashion in the event that the node currently hosting that group fails.

We will leave the clustering subject now and return to the creation of the infrastructure to support our clusters.

Forest Creation Process

Assuming we are starting from scratch, a so-called green fields site, we must first create a forest into which your systems will be integrated. This is called the forest creation process. This is the process that starts with the provisioning of an installation server through the creation of the forest.

Installation of Support Server

The first server installed in your network, you may be surprised to know, is not a domain controller. It is not even a new server. It should be a non-service server installed with either Windows 2000 or Windows Server 2003 in its own workgroup. This server is placed on the lab or future production subnet, initially as a workgroup server, and exposes a number of shares used for accessing operating systems, tools, software, utilities, and patches. The idea is to provide a secure, closed network that does not have access to the outside network that might likely contaminate your implementation. The support server is used for patches, access to tools, resource kits, and so on.

It is critical at this stage that none of your new servers "touch" the Internet or are exposed to the outside. It is very easy to "catch" a virus and not notice it until the entire forest is created and all your servers start croaking.

This server is eventually joined to the network as a temporary *Windows Update Server* (WUS). The server may also function as a temporary DHCP server. To configure the support server, do as follows:

1. Log on to support server as Administrator while this server is still in the lab.
2. Create a folder named C:\ADSTUFF and share as ADSTUFF (actually any name will do).
3. Create a folder named C:\ADSTUFF\Adminpak\.
4. Create a folder named C:\ADSTUFF\Support\.
5. Create a folder named C:\ADSTUFF\Exchange Tools\.
6. Create a folder named C:\ADSTUFF\SQL Server Tools\.
7. Create a folder named C:\ADSTUFF\QA documents\.
8. Crate a folder named C:\ADSTUFF\Scripts\.
9. Create a folder named C:\ADSTUFF\RKTools\.
10. Copy needed tools, MSI files, scripts, data, packages, and so on to these folders.
11. Install anti-virus services and make sure the support server has the latest anti-virus DAT files and is performing the correct scans of its file system.
12. Install Software Update Services Software Update Services on the support server.

13. If needed, create distribution folders for operating system images. You can call the shares STDINST for the Windows Server 2003 Standard Edition or ENTINST for the Windows Server 2003 Enterprise Edition operating system.
14. If needed, create the distribution folders named C:\WEBINST and share as WEBINST for the Windows Server 2003 Web Edition operating system.
15. If needed, create the distribution folders named C:\XPINST and share as XPINST for the XP workstation images.
16. Create distribution shares (for example, C:\..\I386) and copy installation sub-folders and files to the distribution shares (see Table 6.2). This process can be done automatically using the Setup manager utility (setupmgr.exe) on the operating system CD's Support, Tools folder. Setupmgr is found in the deploy.cab file.
17. Configure Software Update Services on the installation.
18. Validate this server (including last scan for anti-virus).

Table 6.2 Example Configuration of Support Server

Item	Description
Server Name	SHQPSERVER
Server IP address	10.10.20.6
W2K3 STD install share	STDINST
W2K3 ENT install share	ENTINST
AD/W2K3 Tools	\.\ADSTUFF
Server Administrator account	Administrator (local machine)
Password	(see your specs)

With the support server in place on your isolated network, you can begin working on the creation of the forest and the domains, accessing your server for support materials as if it were your own mini Microsoft.com Web site.

Installation

Upon installation of the support server to the isolated network, proceed to the installation procedures.

1. Rack and stack your servers in the production racks or on the data center floor with access to the isolated network.
2. Power up the support/installation server.
3. Log on to the installation server as Administrator on the isolated subnet.
4. Reset the Administrator password.
5. Change the IP configuration to statically assigned addressing.
6. Assign the IP address of 10.10.20.23 (on a /22 subnet where the 10.10.20.0 space is reserved for the data center servers).
7. Assign the same IP address as the gateway and DNS.
8. Install DHCP and configure the new scope for the subnet (see Table 6.3).

At this point, the installation and provisioning of the support server is complete.

Table 6.3 Example Configuration of DHCP on the Support Server

Scope	Split	Excluded Range	Default Gateway	CIDR	WINS	DDNS
10.10.20.0 to 10.10.20.254	NA	10.10.20.1 to 10.10.20.20	NA	/22	10.10.20.23	10.10.20.23

After the DHCP server has been installed on the support server, reserve IP addresses for the root domain controllers. This ensures the root DCs obtain the correct IP addresses as soon as their IP configuration is changed from static addressing to DHCP-assigned.

One note to consider before we move on: The subnet we have used here will provide sufficient addresses to meet the needs of a high availabilty network. Don't short change yourself with IP addresses. You should be good to go with a subnet that provides more than a thousand

IP addresses. High availability systems use a lot of IP addresses. A typical two-node cluster should be allocated a block of about 24 addresses for future expansion.

Installation of Root Domain

This section covers the promotion of the root domain controllers. By promoting root domain controllers, we are, in fact, creating the forest in which all future high-availability systems will be installed (see Chapter 5 for the discussion of the dual domain [root-and-child] model). The prerequisite to this process is installation of the operating system (Windows Server 2003, Standard Edition) to the domain controller computers on a RAID-1 array. See Chapter 4 for instructions on the configuration of RAID-1 on this server. The servers should be configured for second and third RAID-5 arrays as required.

It is critical this process completes and proceeds as described herein. Deviation from the process or shortcuts may render the root domain useless and it will have to be rebuilt. The updating of the domain controller servers with the required software updates and security patches can take place after promotion, QA, and validation. (See Chapter 5 for the overall architecture this implementation supports.)

Process

Name the Root Domain DCs. Upon completion of the server installations, the root domain controllers will be given miscellaneous names, and they will be a member of the workgroup setup on the support server. Change the names of the root domain controllers to the names provided in your Active Directory Architecture (discussed in Chapter 5). For the corporate hub, the server names we use here are HQRDC01 and HQRDC02 (for later implementation).

It is important to remember to rename the servers to their DC names prior to running DC promo. The names cannot be changed after promotion of these servers to domain controllers, and they have to be destroyed if the names are incorrect. Do not change the workgroup when changing the names.

Configure TCP/IP on HQRDC01. Log on as Administrator to the server designated to become the root DC (HQRDC01). Open the

TCP/IP properties of the network interface card (NIC), and enter the parameters listed in Table 6.4.

Table 6.4 TCP/IP Configuration on HQRDC01 Support Server

Resources (RDC01)	Configuration
IP	10.10.20.21
Subnet Mask	255.255.252.0
Default Gateway	10.10.20.1
Preferred DNS	10.10.20.21
Alternate	<null>

Configure TCP/IP on HQRDC02. Log on as Administrator to the server designated to become the root DC (RDC02). Open the TCP/IP properties of the NIC, and enter the parameters listed in Table 6.5.

Table 6.5 TCP/IP Configuration on UVRDC02 Support Server

Resources (HQRDC02)	Configuration
IP	10.10.20.24
Subnet Mask	255.255.252.0
Default Gateway	10.10.20.1
Preferred DNS	10.10.20.24
Alternate	<null>

To install DNS, do as follows:

1. Log on as Administrator to the server designated to become the root DC (HQRDC01) and install DNS on this server. This is achieved by opening Control Panel, Add or Remove Programs, and Add/Remove Windows Components. This launches the Windows Components Wizard.

2. Select Networking Services in the wizard and click the Details button. In the Networking Services dialog box, check the option to install Domain Name System (DNS).
3. Complete the procedures and, when prompted by the installation procedure for the Windows Server operating system CD, provide a CD or browse to the I386 folder under the STDINST share (the source for OS installation files) on the installation or support server.
4. Complete the process to install DNS on the server. Repeat the process for all hub root domain controllers.

Now you can create the Forest Root Zone on HQRDC01. To create the forest root zone, perform the following steps (note: this process is not repeated on HQRDC02 or any other root server destined to become a DC):

1. Start DNS and right-click on the HQRDC01 icon.
2. Select New Zone. The New Zone Wizard launches. Click Next.
3. Select the option to create a Primary zone and click Next.
4. Select Forward Lookup zone and click Next.
5. Enter the domain name (such as MCITY.CTY) as the name of the zone and click Next.
6. Keep the default DNS file name (it should be MCITY.CTY.dns) for the zone file name and click Next.
7. If prompted for Dynamic Update configuration, choose the option to allow Dynamic Updates. Click Next.
8. Complete the process by selecting Finish.

Create the Reverse Lookup Zone on HQRDC01. To create the reverse lookup zone for the forest, perform the following steps:

1. Open the DNS console and expand the HQRDC01 server icon.
2. Select Reverse Lookup Zones and click on New Zone. The New Zone Wizard launches.
3. Select options for a Primary non-integrated zone and click Next.
4. Enter the IP address range for the zone; this is the 10.10.20.X network.
5. Click Next and select the options to enable dynamic update.
6. Complete the process by selecting Finish.

Create the Forest Root Domain Controller on HQRDC01. To create the forest root domain, perform the following steps:

1. Click Start, Run, and type **DCPROMO** on HQRDC01.
2. Choose the options for creating a root domain controller in a new forest.
3. Choose the root domain name as the full DNS name for the new domain (MCITY.CTY).
4. Accept the default NetBIOS name for the domain.
5. Choose the default path for the SYSVOL folder on the RAID-5 array. However, the drive letter should point to the RAID-5 array on (D, E, or F) and not C:\ (for example E:\Windows\…). Choose the path options provided for the NTDS Active Directory database and its log files, changing only the drive letters to point to the RAID 5.
6. Accept permissions compatible with Windows 2000 and Windows Server 2003.
7. Enter the Directory Services Restore Mode Administrator password (this should be a complex password, choose something like 4NTDS@mcity), ignoring the quotes. (Remember the server's local Administrator password becomes the password required to log on to the DC after promotion.)

Review the settings, and click Finish to begin the process. Restart the server when prompted.

Enable Active Directory Integration of the Forest Root Zone and the Reverse Lookup Zone. To enable AD integration for the root zone, do as follows:

1. Open the DNS console and expand the root server HQRDC01 icon.
2. Expand the Forward Lookup Zones folder and select the MCITY.CTY zone. Right-click this zone and select Properties.
3. The Properties dialog box for MCITY opens. On the General tab, select the Change button on the Type option. The Change Zone Type dialog box launches.
4. Select the option to change the zone to Active Directory Integrated and click OK.

Perform the same procedure on the Reverse Lookup Zone folder. Verify HQRDC01 Name Registration. To verify name registration, perform the following actions:

1. Open the DNS console and expand the root server HQRDC01 icon.
2. Expand the Forward Lookup Zones folder and select the MCITY.CTY zone.
3. Verify whether _msdcs, _sites, _tcp, and _udp sub-domains are registered under MCITY.CTY.
4. If these sub-domains are not registered, then start a command prompt and type **NET STOP NETLOGON**. Wait for the service to stop and then type **NET START NETLOGON**.
5. Repeat steps 1 through 3 to verify the registration.
6. Verify the Reverse Lookup Zone has replicated.

Verify DNS name resolution on HQRDC02. Before HQRDC02 can be promoted as a root DC, DNS first must be verified. This can be achieved as follows:

1. Log on to HQRDC02 as the Administrator.
2. Open the command prompt and type **NSLOOKUP MCITY.CTY** and press Enter. You should see the following result:

```
C:\>nslookup MCITY.CTY
Server: HQRDC01.MCITY.CTY
Address: 10.10.20.21
Name: MCITY.CTY
Address: 10.10.20.21
```

If you do not see this, check to see whether the IP settings on HQRDC02 are correct. It should have HQRDC01 (10.10.20.21) as its preferred DNS server. Do not proceed with DCPROMO of HQRDC02 until DNS is working properly.

Perform DCPROMO on the server HQRDC02. To create the second domain controller, perform the following steps:

1. Click Start, Run, and type **DCPROMO** on HQRDC02.
2. Choose the options for creating an additional domain controller for an existing domain and click Next.

3. You are prompted for access to the root domain. Choose the Administrator account because this account has Enterprise Administrator credentials. See the previous steps for account and password information.
4. Choose the default path for the SYSVOL folder on the RAID-5 array. However, the drive letter should point to the RAID-5 array on (D, E, or F) and not C:\. Choose the path options provided for the NTDS Active Directory database and its log files, changing only the drive letters to point to the RAID 5 volume as previously mentioned (see Chapter 4).
5. Enter the Directory Services Restore Mode Administrator password for this server (this should be a complex password; choose 4NTDS@MCITY). DCs can and should have the same Directory Services Restore Mode Administrator password to simplify administration.

Review the settings and then click Finish to begin the process. Restart the server when prompted. Verify HQRDC02 Name Registration. To verify name registration, perform the following actions:

1. Open the DNS console and expand the root server HQRDC02 icon.
2. Expand the Forward Lookup Zones folder and select the MCITY.CTY zone.
3. Verify whether _msdcs, _sites, _tcp, and _udp sub-domains are registered under MCITY.CTY.
4. If the sub-domains are not registered, then start a command prompt and type **NET STOP NETLOGON**. Wait for the service to stop and then type **NET START NETLOGON**.
5. Repeat steps 1 through 3 to verify the registration.

Verify the Reverse Lookup Zone has replicated. Update the Preferred DNS Parameters on HQRDC01. Log on to HQRDC01 and open the TCP/IP properties for the NIC. Change the preferred DNS server from 10.10.20.21 to 10.10.20.24.

Create *Automated System Recovery* (ASR) media for the domain controllers. The creation of the root domain and promotion of the first domain controllers is now complete. System recovery using ASR media now must be performed on the domain controllers. After the ASR disks have been created, you can start the QA discussed in the next section.

Quality Assurance

QA and validation must be performed before continuing further. QA can be achieved by following these steps:

1. Join a clean Windows XP SP1, SP2, or higher workstation to the root domain. Remember to follow the naming convention for the workstation according to Active Directory Architecture.
2. Install the WSO3 support tools on the workstation. The tools can be accessed from the ADSTUFF\SHQPORT\TOOLS share on the installation server. Install the tools to the default path on the C: drive.
3. Install the ADMINPAK on the workstation. This installs management tools, such as DSA.MSC, to the workstation. The tools can be accessed from the ADSTUFF\ADMINPAK share on the installation server. Install the tools to the default path on the C: drive.
4. Install the Resource Kit tools to the workstation. This installs tools, such as DNSDIAG, DCDIAG, and DSQUERY to the workstation. The tools can be accessed from the ADSTUFF\RESKIT share on the installation server. Install the tools to the default path on the C: drive.
5. Open a command console and run DCDIAG /s:<domain controller name> /a /f<logfile> /ferr<errlogfile>. Perform the DCDIAG against both HQRDC01 and HQRDC02. The data generated by DCDIAG is piped to the default log file location on the workstation.
6. Perform DCDIAG several times a day during the installation.
7. Open the replication monitor and check that replication is occurring without errors between the domain controllers.

Finally, you can run DSQUERY against the domain controllers to see that all FSMO roles are intact (the roles are moved later on in the implementation). Much of this manual diagnostics and QA can be left to Microsoft Operations Manager (MOM) to handle. Without MOM, QA can become something of an endurance during the life of a long project to stand-up a high availability infrastructure.

Forest Preparation, DNS, and Exchange

This section covers the preparation of the forest for the addition of Exchange 2003 and the creation of a child domain. Before the child domain can be created, the schema is extended in the forest root to cater to the addition of domains and the installation of Exchange 2003 into the forest. All activities in the root domain on the domain controllers are done using the root domain's Administrator account. In addition, we prepare the first site and associate it with the new subnet and perform some housekeeping. The process of preparing the forest is outlined as follows.

Move Domain Operations Master Roles. HQRDC01 is a GC server and also holds schema and domain naming operations. It is important to move the domain operations roles to HQRDC02.

1. Start Active Directory User and Computers on HQRDC01.
2. Right-click the root node and select Connect to Domain Controller. Choose HQRDC02 and click OK. You are now on RDC02.
3. Right-click the MCITY.CTY domain and select Operations Masters.
4. The RID Master Role appears first. Select Change to move it to HQRDC02. You are able to select the target computer from a list if necessary.
5. Click Yes to confirm the transfer.

Repeat these steps for the PDC Emulator and Infrastructure Master roles.

Configure DNS Forwarders. We cannot leave the DNS Servers in the root domain as root servers because they assume they have root authority and users are unable to resolve addresses in the external namespaces the county owns and on the Internet. To add DNS forwarding, we first have to delete the root DNS zone (if it exists) and add DNS forwarder addresses. This is done as follows:

1. Right-click the "." folder under the Forward Lookup Zones and select Delete.
2. Right-click the DNS Server name HQRDC01 and select Refresh.

3. Right-click the server name again and select Properties. Click the Forwarders tab and check the Enable Forwarders check box.
4. Enter the IP addresses of the external, private DNS server (primary and alternate) for MCITY.CTY (most likely these are your external domain's ISP) and click Add. (The addresses of these servers are usually obtained from the Datacom or Network Group in your enterprise.)

When all addresses are entered, click OK to close the dialog box.

Perform ASR backups of the root domains controllers.

Verify Credentials of the root domain's Administrator Account. Before continuing with schema changes, it is worthwhile to confirm the credentials of the Administrator because, at this stage, no other account has the rights needed to perform forest operations:

1. Open Active Directory User and Computers (DSA.MSC) and expand the Users folder.
2. Double-click the Enterprise Admins group and verify whether the Administrator account is present. Add the account if it is not.
3. Perform the same verification on the Schema Admins group. If Administrator is not present, then add the account. Close down the DSA.
4. Prepare the forest for Exchange 2003. This process requires the Exchange 2003 Installation CD.
5. Insert the CD into the drive. If the Exchange installation process boots, then close it down.
6. Open the command prompt on HQRDC01 and enter the following command: **<CDDRIVE>:\setup.exe /ForestPrep**. During ForestPrep, you are prompted for the account for the Full Exchange Administrator. Use only the MCITY\Administrator account (which can be removed later).
7. When the ForestPrep completes, remove the Exchange CD and check for any errors in the event logs. Report any errors related to ForestPrep for review and copy the ForestPrep progress log files to the Service Admin workstation. (It is important that the DC on which you perform this can "see" the DC that has the Schema Master role, otherwise ForestPrep will fail.)
8. Allow root domain controllers enough time to replicate the changes made before moving onto the next step. Never rush into the next step, and if possible give your new domain at least 12

hours before continuing. You can check to see whether the schema additions have replicated by confirming the presence of the exchange object in ADSI Edit. Connect ADSI Edit to both domain controllers. They should both show the exchange objects in the Configuration container.

9. Review DCDIAG results during the replication cycle as previously described.

Change the Default-First-Site-Name. This step takes place after all forest changes have replicated and DCDIAG results are normal. To do this, do the following:

1. Start Active Directory Sites and Services.
2. Expand Sites.
3. Right-click Default-First-Site-Name and select Rename. Enter HQ (corporate headquarters).
4. Do not close down the console.
5. Add the Subnet Associated with the Site. In this step, we add the subnet and mask associated with HQ.
6. With the console still open, right-click the Subnets folder and select New Subnet.
7. Enter the Network ID (10.10.20.0) for the subnet associated with HQ (you need to select HQ in the sites list).
8. Click OK and close the console.
9. Perform QA.
10. Create new ASR media.

You are now ready to move on to a child domain.

Installation of Bridgehead Servers and the Child Domain

This section outlines the steps required to build the bridgehead servers for the main hub site (HQ) and promote the domain controllers into the child domain AD.MCITY.CTY. The process is outlined in the next section. You typically do not need to specifically set up bridgehead servers on small domains (with less than 100 domain controllers).

The first procedure to perform on the bridgehead or sub-domain controllers is the configuration of DNS, particularly forwarding. The process is similar to the configuration of DNS in the root domain.

1. **Name the Child Domain DCs:** Upon completion of the server installations, the child domain controllers are given miscellaneous names, and they are a member of the workgroup created when you installed the support server. Change the names of the child domain controllers to the names provided in your Active Directory Architecture. For the HQ hub, the server names are HQSDC01 and HQSDC02. It is important to remember to rename the servers to their DC names prior to running DC promo. The names cannot be changed after promotion of these servers to domain controllers, and they have to be destroyed if the names are incorrect. Do not change the workgroup when changing the names.
2. **Configure TCP/IP on HQSDC01:** Log on to the server designated to be promoted first (HQSDC01) as Administrator. Open the TCP/IP properties of the NIC and enter the parameters listed in Table 6.6.

Table 6.6 TCP/IP Configuration of SDC01 Domain Controller

Resources (SDC01)	Configuration
IP	10.10.20.27
Subnet Mask	255.255.252.0
Default Gateway	10.10.20.1
Preferred DNS	10.10.20.27
Alternate	10.10.20.30

3. **Configure TCP/IP on HQSDC02:** Log on to the server designated to become the second DC to be promoted (HQSDC01) as Administrator. Open the TCP/IP properties of the NIC and enter the parameters listed in Table 6.7.

Table 6.7 TCP/IP Configuration of SDC02 Domain Controller

Resources (SDC01)	Configuration
IP	10.10.20.30
Subnet Mask	255.255.252.0
Default Gateway	10.10.20.1
Preferred DNS	10.10.20.30
Alternate	10.10.20.27

4. **Install DNS:** Log on as Administrator to the server designated to be promoted first (HQSDC01), and install DNS on this server. This is achieved by opening Control Panel, Add or Remove Programs, and Add/Remove Windows Components; this launches the Windows Components Wizard. Select Networking Services in the wizard and click the Details button. In the Networking Services Dialog box, check the option to install DNS. Complete the procedures and, when prompted by the installation procedure for the Windows Server operating system CD, provide a CD or browse to the I386 folder under the STDINST share on the installation server.
5. **Complete the install:** Finish the process of installing DNS on the server. Repeat the process for all hub child domain controllers (prior to promotion).

Configure Forwarding. Log on as Administrator to the server designated to be promoted first domain controller (HQSDC01), and open the DNS console to configure forwarding:

1. Right-click the server name and select Properties.
2. Select the Forwarders tab.
3. Check the option Enable Forwarders.
4. Enter the IP addresses for the forest root servers, HQRDC01, HQRDC02, and DRRDC01 (the root domain controller at the DR site).
5. Check the option Do Not Use Recursion.
6. Click Apply and close the DNS console.

7. Verify forwarding by performing a NSLOOKUP on AD.MCITY.CTY. The lookup should fail because the zone for the AD domain is not yet created. However, you should see whether the query correctly forwarded to one of the root DCs. If you get a timeout on the request, then forwarding is not set up correctly. Check the forwarder settings and, if you are still getting a timeout, go down to the network layer and make sure the child DCs can ping the root DCs.

Repeat these steps on all child domain controllers.

Delegate the AD DNS domain to HQSDC01. To delegate the child domain, you need to open the DNS console on the root DC HQRDC01. This can be achieved by opening the DNS console on the service admin workstation that has an account in the root domain:

1. Select the MCITY.CTY domain and right-click. Select New Delegation. The New Delegation Wizard launches. Click Next.
2. In the Delegated domain field, enter the name of the domain to be delegated, namely AS (the FQDN is, thus, AD.MCITY.CTY). Click Next.
3. On the Name Servers page, click Add. This lets you enter the name of the DNS server that hosts the sub-domain. Enter the FQDN of the server in the server FQDN field and attempt to resolve the IP address of the server. If you cannot resolve the name (which is likely to be the case at this point), then enter the known IP address for the HQSDC01 server. Click OK and then click Next.
4. Click Finish to create the delegation.

Create the DNS zone on HQSDC01. To create the primary DNS zone for AD, perform the following steps (this process is *not* repeated on HQSDC02 or any other server destined to become a DC in the HQ hub):

1. Start DNS and right-click the HQSDC01 icon.
2. Select New Zone. The New Zone Wizard launches. Click Next.
3. Select the option to create a standard Primary zone and click Next.
4. Select Forward Lookup zone and click Next.
5. Enter **AD.MCITY.CTY** as the name of the zone and click Next.

6. Keep the default DNS file name (it should be AD.MCITY.CTY.dns) for the zone file name, and click Next.
7. If prompted for Dynamic Update configuration, choose the option to allow dynamic updates. Click Next.
8. Complete the process by selecting Finish.

Create the Child Domain Controller and Domain on HQSDC01. To create the child domain, perform the following steps:

1. Click Start, Run, and type **DCPROMO** on HQSDC01.
2. Choose the options for creating a domain controller for a new domain; that is, select the domain controller for a new domain option. Click Next.
3. Select the option Create a New Child Domain in an Existing Domain Tree, and then click Next.
4. Enter the Enterprise Administrator credentials (MCITY\Administrator), and in the Domain box, enter **AD.MCITY.CTY**.
5. Provide MCITY.CTY as the parent domain.
6. Enter AD as the child domain and click Next.
7. Click Next to accept the default NetBIOS name AD.
8. Choose the default path for the SYSVOL folder on the RAID-5 array. However, the drive letter should point to the RAID-5 array on (D, E, or F) and not C:\ (for example, E:\Windows\…). Choose the path options provided for the NTDS Active Directory database and its log files, changing only the drive letters to point to the RAID 5 volume as previously mentioned (see Chapter 4).
9. Click OK if you receive a message indicating the DNS server for the domain was not found. This occurs if there is no A record for the domain yet.
10. Accept permissions compatible with Windows 2000 and Windows Server 2003.
11. Enter the Directory Services Restore Mode Administrator password (this should be a complex password, so choose something like 4NTDS@ MCITY). Remember the server's local Administrator password becomes the password required to log on to the DC after promotion.

Review the settings and click Finish to begin the process. Restart the server when prompted. Enable Active Directory Integration of the AD Zone. To enable Active Directory integration for the zone, do as follows:

1. Open the DNS console and expand the root server HQSDC01 icon.
2. Expand the Forward Lookup Zones folder and select the HQ.MCITY.CTY zone. Right-click this zone and select Properties.
3. The Properties dialog box for AD opens. On the General tab, select the Change button on the Type option. The Change Zone Type dialog box launches.
4. Select the option to change the zone to Active Directory Integrated and click OK.

Verify HQSDC01 Name Registration. To verify name registration, perform the following actions:

1. Open the DNS console and expand the root server HQSDC01 icon.
2. Expand the Forward Lookup Zones folder and select the AD.MCITY.CTY zone.
3. Verify whether _msdcs, _sites, _tcp, and _udp sub-domains are registered under AD.MCITY.CTY.
4. If these sub-domains are not registered, then start a command prompt and type **NET STOP NETLOGON**. Wait for the service to stop and then type **NET START NETLOGON**.
5. Repeat steps 1 through 3 to verify the registration.

Verify DNS name resolution on HQRDC02. Before HQSDC02 can be promoted as an additional child DC, DNS first must be verified. This can be achieved as follows:

1. Log on to HQSDC02 as the Administrator.
2. Open the command prompt and type **NSLOOKUP AD.MCITY.CTY** and press Enter. You should see a resolution to AD.MCITY.CTY from 10.10.20.27.

If you are not able to resolve the domain, check to see whether the IP settings on HQSDC02 are correct. It should have 10.10.20.30 as its preferred DNS server address and 10.10.20.27 as the alternate. Do not

proceed with the DCPROMO of HQSDC02 until DNS is working properly. DCPROMO the HQSDC02 server. To create the second domain controller, perform the following steps:

1. Click Start, Run, and type **DCPROMO** on HQSDC02.
2. Choose the options for creating an additional domain controller for an existing domain and click Next.
3. You are prompted for access to the child domain. Choose the Administrator account for AD. The Administrator password is the same as the server password before the DC was promoted.
4. Choose the default path for the SYSVOL folder on the RAID-5 array. However, the drive letter should point to the RAID-5 array on (D, E, or F) and not C:\. Choose the path options provided for the NTDS Active Directory database and its log files, changing only the drive letters to point to the RAID 5 volume as previously mentioned (see Chapter 4).
5. Enter the Directory Services Restore Mode Administrator password for this server (this should be a complex password; choose 4NTDS@MCITY). DCs can and should have the same Directory Services Restore Mode Administrator password to simplify administration.

Review the settings and then click Finish to begin the process. Restart the server when prompted. Verify HQSDC02 Name Registration. To verify name registration, perform the following actions:

1. Open the DNS console and expand the root server HQSDC02 icon.
2. Expand the Forward Lookup Zones folder and select the AD.MCITY.CTY zone.
3. Verify whether _msdcs, _sites, _tcp, and _udp sub-domains are registered under MCITY.CTY.
4. If these sub-domains are not registered, then start a command prompt and type **NET STOP NETLOGON**. Wait for the service to stop, and then type **NET START NETLOGON**.
5. Repeat steps 1 through 3 to verify the registration.

Move the Domain Operations Master Roles. HQSDC01 is a GC server and becomes the preferred bridgehead server (Active Directory does this automatically in Windows Server 2003). To lessen the load on this

server, the RID operations master needs to be moved to HQSDC02. We also should move the IF and PDC roles to HQSDC02.

1. Start Active Directory User and Computers on HQSDC01.
2. Right-click the root node and select Connect to Domain Controller. Choose HQSDC02 and click OK. You are now on HQRSC02.
3. Right-click the AD.MCITY.CTY domain and select Operations Masters.
4. The RID Master Role appears first. Select Change to move it to HQSDC02. You are able to select the target computer from a list if necessary.
5. Click Yes to confirm the transfer.

Repeat these steps for the PDC Emulator and IF roles. Create ASR media for the Domain Controllers. The creation of the child domain and its controllers is now complete. System recovery using ASR media now must be performed on the domain controllers.

Quality Assurance and validation must be performed before continuing further. QA can be achieved by following these steps:

1. Open a command console and run DCDIAG /s:<domain controller name> /a /f<logfile> /ferr<errlogfile>. Perform the DCDIAG against both HQSDC01 and HQSDC02. The data generated by DCDIAG is piped to the default log file location on the workstation.
2. Perform DCDIAG several times a day and run DCDIAG at the enterprise level.
3. Open the replication monitor REPLMON, and check to see whether replication is occurring without errors between the domain controllers. You also can use REPADMIN to check the update vectors and force replication between the replication partners.
4. Finally, you can run DSQUERY against the domain controllers to see that all FSMO roles are intact.

Also, load the replication monitor and ensure infrastructure changes between the domains are replicating. (See Chapter 13 on using MOM for alerts and monitoring.) The Windows Server 2003 replication service (FRS) Management Pack can get this going for your in no time flat.

Installing DHCP and WINS Services

This section covers installing the DHCP and WINS services on the sub-domain controllers. These services are essential to support the client network so they do not have any problem accessing the client network. In this process, we first install the services for both DHCP and WINS, and then we configure them to provide the required services. For the time being, we only configure the DHCP servers to service the 10.10.20.0 subnet and import the scopes from the legacy DHCP servers (should we need to) after the implementation is released to production.

The process of installing the DHCP and WINS service is as follows:

1. While keeping the standby DHCP server going, provide a static IP address for HQSDC01 and HQSDC02. These addresses can be the ones that were originally assigned these servers during the creation of the sub-domain. Notice that we are not going to install DHCP on the root domain controllers. Ensure that these addresses cannot be obtained by other clients on the network. This is achieved by reserving a range of IP addresses for static assignment.
2. Install the DHCP and WINS service from the Windows Components facility on both HQSDC01 and HQSDC02. The files required for the installation can be obtained from the \STDINST\i386 directory on the installation server we discussed at the beginning of this chapter.
3. Upon completion of the installation of the DHCP servers, authorize both servers in Active Directory. This can be done in the root domain using the DHCP console. Log on to the root domain as the Administrator. This operation requires Enterprise Administrator credentials, which the root Administrator has.
4. On HQSDC01 only, create a scope called "Server Subnet" with the parameters listed in Table 6.8. Upon creation of the scope, ensure the scope is deactivated (the default upon creation is deactivated). To deactivate the scope, right-click the scope name and select Deactivate.

Table 6.8 DHCP Server Settings

Parameter	Value
Scope Name	Server subnet (for example, 10.10.20.0)
Scope	10.10.20.1 to 10.10.22.254
Exclusion	10.10.20.1 to 10.10.20.41
003 Router	10.10.20.1
015 DNS Server	10.10.20.27, 10.10.20.30
044 WINS/NBNS Servers	10.10.20.27, 10.10.20.30
015 DNS Domain Name	AD.MCITY.CTY
046 WINS/NBT Node Type	0x08

As soon as the scope has been configured according to the values in Table 6.8, you can deactivate and stop the DHCP service on the installation server. After this service is offline, you can activate the scope on HQSDC01. First, however, you should create the superscope as outlined here:

1. To create the superscope, select the server icon (HQSDC01) and right-click.
2. Select New Superscope and click Next.
3. Provide a name for the superscope. You can name it something like HQ on HQSDC01. Click Next (do not add the word "Superscope" to the name because Windows adds it anyway).
4. Add the HQ scopes this server is hosting to the superscope.
5. Click Next and then Finish to create the superscope.

At this point, there are no other scopes to configure on the server. There is also no WINS server configuration required at this time. However, after setting up the DHCP servers, DHCP server HQSDC02 must be configured to offer the legacy scopes that might be migrated from any legacy DHCP servers on a preexisting network. To migrate the scopes (from a Windows 2000 or Server 2003 machine), you need to export the

DHCP settings from the legacy server that is being replaced by HQSDC02. This can be achieved as follows:

1. Run DhcpExim.exe (this tool is available on the CD that accompanies the Windows Server 2003 Deployment Kit; it can also be downloaded from Microsoft) against the DHCP server from which you are exporting.
2. In the DhcpExim dialog box that loads, select the option Export Configuration of the Local Service to a File.
3. In the DHCPEXIM Export to a File dialog box, enter the file name and location to save the file, then click OK.
4. In the DhcpExim Export dialog box, select all the scopes on the list to migrate; ensure you select to migrate all the settings on the server. Do not select the Disable the Selected Scopes on the Local Machine Before Export option because the scopes are needed to continue service for a few weeks after the export.
5. Click OK and wait for the message "The operation was completed successfully."
6. Copy the exported file to media that can be installed onto the installation server on the isolated 10.10.20.0 network. You cannot download the setting through any network connection between the old DHCP server and the new one.
7. To import the settings into HQSDC02 where the second DHCP server is running, open the command console and enter the command **netsh DHCP server import <path to export file> all**.

Confirm the import of the scopes into the HQSDC02 and then place them under the superscope called Legacy for HQ site, creating the superscope as demonstrated earlier. For WINS, ensure the WINS services are installed on both child domain controllers; however, only activate the WINS server on the DC that is not holding the PDC role. Ensure the superscope is deactivated until needed.

Patching and Updating Domain Controllers

Before continuing to configure any further services in the root domain, the root domain controllers should be patched and updated from the makeshift Windows Update Server (WUS). By installing WUS (or its

predecessor *Software Update Services [SUS]*)on the isolated network, all servers in the isolated network can pull down updates and patches to bring them to the latest patch level and to ensure that critical software updates, especially security updates, are applied.

The process of ensuring all root domain servers are regularly updated while being staged on the isolated network is as follows:

1. Create a new group policy object at the domain level (root and child domains).
2. Call the policy SoftwareUpdatePolicy.

This GPO is created so the policy can persist after security templates and group policy are imported into the default domain policy later in the installation process.

After the policy is configured, all Windows 2003 Server, Windows 2000 Server SP3 and later, and all Windows XP clients are serviced by the SUS server.

WUS is beyond the scope of this book, but it is critical in the maintenance of a high availabilty network. Microsoft publishes excellent white papers on update services. They can be accessed on the Microsoft Web site (simply search for SUS or WUS).

Exchange Domain Preparation

This section covers the preparation of the child domain for the addition of Exchange 2003. This process entails running the /DomainPrep switch using Exchange setup.exe, similarly to what was performed during the /ForestPrep process. The domain preparation creates the groups and permissions in Active Directory necessary for Exchange operation. Domain prep also creates several folders and performs a variety of tasks needed to install Exchange 2003 into the domain. Obviously you cannot install an Exchange cluster like we will do in Chapter 10 until this process is complete.

Prepare the domain for Exchange 2003. This process requires the Exchange 2003 Installation CD as is described as follows:

1. Log in as Administrator to the *primary domain controller* (PDC) emulator in the child- or sub-domain. Domain Admin credentials are needed for this process. The account does not need

Enterprise Administrator or Schema Administrator credentials; therefore, in this process, AD\Administrator will do.

2. Insert the CD and, if the Exchange installation process boots, then close it down.
3. Open to the command prompt on the PDC DC and enter the following command: **<CDDRIVE>:\setup\i386\setup.exe /DomainPrep**.
4. When the DomainPrep completes, remove the Exchange CD and check for any errors in the event logs. Report any errors related to the process for review and copy the DomainPrep progress log files to the Service Admin workstation. Reboot the DC if required.
5. Allow root domain controllers to replicate the changes made before moving onto the next step.

Review DCDIAG results during the replication cycle in Step 4. (See also Chapter 13 on deploying Microsoft Operations Manager for monitoring Exchange.)

Creation of Initial Service and Administration Resources

This section covers the creation of the initial service and administration accounts and services. These accounts and resources are created before security and group policy are applied to the domain and forest. The resources are needed to flesh out the core servers and services of the Windows Server 2003 network and Active Directory.

The process of installing these services is outlined in Figure 6.4.

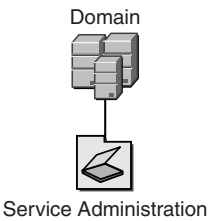

Figure 6.4 Creating Service and Administration Resources.

To create the necessary services, do as follows:

1. Create a Tier 1 OU called Service Administration.
2. Create three child OUs under Service Administration called Users, Computers, and Groups.
3. Select the Users OU, right-click and select New and then select User.
4. In the New Object – User dialog box, enter the first and last name of the user (for example, Mickey Mouse).
5. In the Full Name edit field, follow the guidelines in your Active Directory Architecture for service and administration accounts (make sure the logon account name is the same as the Full Name field).
6. Enter the password provided and then follow the password reset guidelines for the service accounts.
7. Select Finish and create the account.

After the account is created, open the account and provide it the appropriate group membership. Create more of the same account using the Copy account facility (this retains group membership for all accounts). Perform this on the Tier 1 OU called Service Accounts.

Next, create a root OU at the same level as Service Administration, and call it Servers. Under this OU, create three sub-OUs for cluster servers and call them Exchange, SQL Server, and File and Print. The OU namespace is demonstrated in Figure 6.5. The reason for this namespace is such that you can create GPOs for the Exchange, SQL Server, and File and Print Servers, respectively.

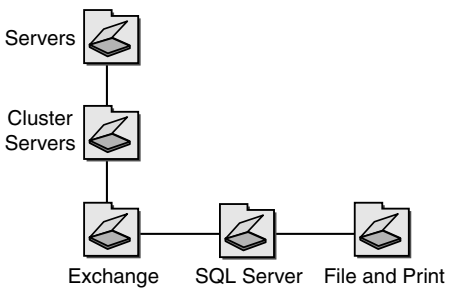

Figure 6.5 Creating OUs for cluster servers.

Next, we install the line of business servers that are configured and clustered for high availability.

Clustering

A number of steps must be completed before a cluster with multiple nodes is complete. Figure 6.6 provides a flow-chart of the steps to peform.

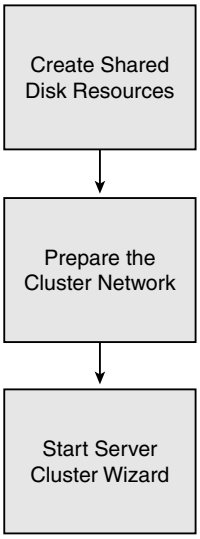

Figure 6.6 Cluster creation flow-chart.

Create Shared Disk Resources

The first step in creating the cluster is the configuration of the disk drives. Obviously, the cluster creation fails if it does not find or recognize drives. If you are going to cluster on a SAN or a SCSI-shared storage array, then you first need to install your *host bus adapters* (HBAs) in the servers and configure them. This step might entail installing special drivers for the cards, management software, and any patches that may be necessary to get them working in Windows Server 2003.

After the adapters are installed and the interface management software sees the controllers working, you'll connect them to the SCSI array or to the switches of the SAN fabric. By now, your disk arrays are installed and ready to go.

This looks like a small step from the flow-chart in Figure 6.6, but it's not. It can take a lot of time and effort to set up the SAN devices, and the effort can vary greatly between different SANs SCSI arrays or disk replication solutions, such as the one provided by NSI-Software (see Chapters 9 and 10). The installation and configuration of the SAN, fabric, zoning, and so on, is very complex and beyond the scope of this book.

Make sure your servers see the external or replicated drives. If everything is configured properly, your Windows servers see the drives as if they are installed on the same server. You are able to manage the new drives the system sees from the SAN management software and various server utilities, including the Computer Management utility. This is demonstrated in Figure 6.7.

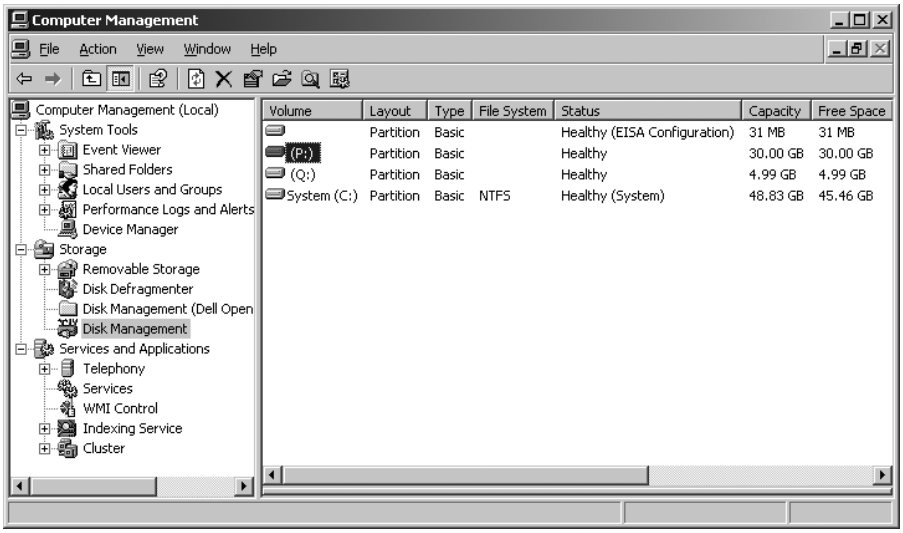

Figure 6.7 Computer Management recognizing the shared disk drives.

In Figure 6.7, notice the presence of the P and Q drives. In this case, we have configured the P drive for the application and the Q drive is the drive that holds the quorum resource. The "Quorum" drive is essential for the cluster and accommodates the so-called Quorum resource.

Prepare the Cluster Network

The next step is to prepare the cluster network or interconnect between the nodes in the cluster. If you are installing a 2-node (active-passive) cluster, it is possible to install an interconnect network between the nodes using a single network cable attached NIC to NIC. The network link needs to be crossed over, but you may not need a cross-over cable because most modern servers employ NICs that recognize the need to cross over the datapaths.

Your interconnect IP configuration must be different to the LAN NICs. In other words, you should set up a private subnet between the servers (unless you are setting up geo-clusters and don't have enough cable to stretch your cluster from NY to LA). For example, if your LAN is on a subnet configured as 10.10.20.0, then put the interconnect on a 192.168.0.0 subnet. The IP on a one-node is, thus, 192.168.0.1, and the NIC on the other node is 192.168.0.2. Leave the gateway addresses on both NICs vacant. As long as the .1 can ping .2, your interconnect is ready.

If you are going to install an N+1 node or any configuration comprising of more than two nodes, then you need to use a hub for your interconnect network. This issue was discussed in Chapter 4. Remember, you don't need a switch.

One last word: Make sure your interconnect NIC's IP addresses do not end up in the DNS configuration as belonging to your virtual server (the cluster name) because that can result in problems for clients connecting to the name resource. In other words, they can look up the resource IP address, but they are unable to connect to it.

Start Server Cluster Wizard

You can install a cluster interactively using the GUI of the Server Cluster Wizard, or from the command line with command-line parameters passed to the "cluster" executable (cluster/create). We recommend that until you know enough about what makes the cluster service tick, you should work with the wizard. The remainder of this chapter discusses installing a cluster using the wizard.

At this point in the cluster configuration and installation, shut down all potential cluster nodes except the first node. It is important you install the first node without the possibility that other nodes might interfere with the installation process. The cluster is created on the first node

because it is allowed to gain exclusive use of the shared resources. It installs a cluster only if it discovers that it is the first node in the cluster. After the cluster has been created, the next node is added to the cluster, and the procedure is different.

Also, when you power on and start the operating system, make sure it is only the first node that has access to the cluster disk. If another server can see and access the disks, the data on them can be easily destroyed and would have to be reformatted. To prevent the corruption of the cluster disks, you should shut down all but the cluster node you are going to make the first node in the cluster. You can use other techniques (such as, *Logical Unit Number* or LUN masking, selective presentation, or zoning) to protect the cluster disks before creating the cluster, but we have learned that it's safer to simply power down all the other nodes until you have a cluster. After the Cluster service is running properly on one node, the other nodes can be powered up and then added to the cluster as needed.

When you create a cluster, the physical disk resources are automatically created for cluster disks that use drive letters. As mentioned earlier, follow a sound naming convention for all your resources and keep the names consistent. This is critical to do as you will see in the final chapter in this book when we configure and use Microsft Operations Manager. The alerts and logs are not much help if you can't identify the devices and the servers they are on in your MOM data.

To get started, open Cluster Administrator from Adminstrative Tools. Select File, Open Connection, and then select Create New Cluster from the Action list in the dialog box that appears. This is demonstrated in Figure 6.8. Click OK to launch the Server Cluster Wizard.

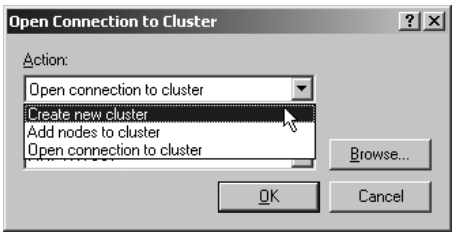

Figure 6.8 The Create New Cluster option in Cluster Administrator.

The first request the wizard makes is for the domain name of the cluster and the cluster name. For the deployment shown in this chapter,

we make sure that we have the correct domain name and that the name you use for the virtual server is the cluster name. This is demonstrated in Figure 6.9. You can set up additional network names for the actual application resoureces (such as SQL Server) as we show in Part II of this book.

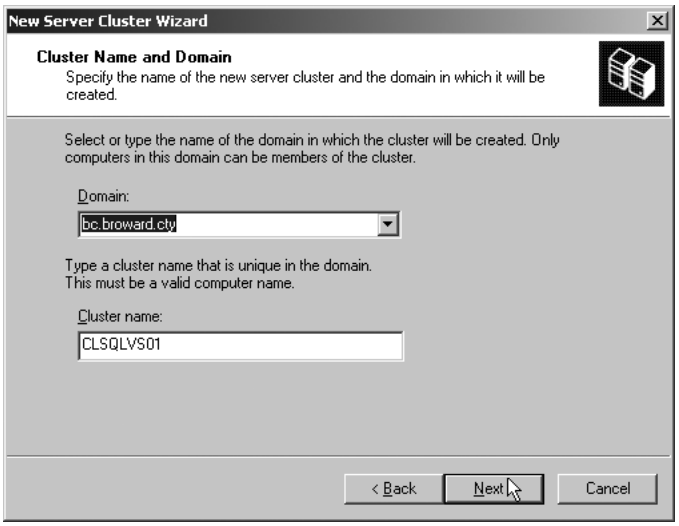

Figure 6.9 Cluster domain and cluster name.

Enter the domain name and cluster name and then click Next. The Select Computer Name dialog box appears. Enter the name of the node you are installing (typically the server on which you started the Cluster Administrator), and then click Next. The wizard now analyzes the configuration to check if it has everything it needs to create a cluster. This is demonstrated in Figure 6.10, which shows the cluster has failed due to a variety of reasons. When you see a lot of red and yellow in the dialog box, it's a sign you have work to do before you can move forward.

If the analaysis fails, you can simply go back or cancel out of the wizard and proceed to fix the problems that were discovered in the configuration analysis stage. The wizard then can be restarted at any time. Figure 6.11 shows that now the configuration analysis has succeeded. You are looking for check marks in all areas and a solid green line on the progress bar. When you have a clean analayis, click Next to continue installing the cluster.

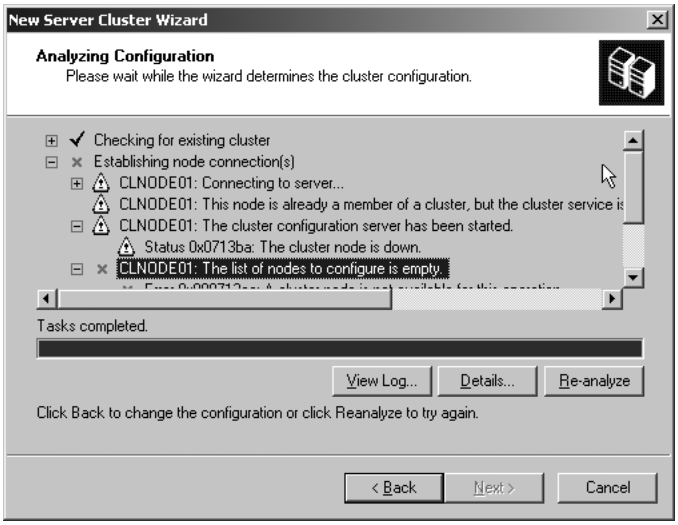

Figure 6.10 Cluster configuration analysis has failed.

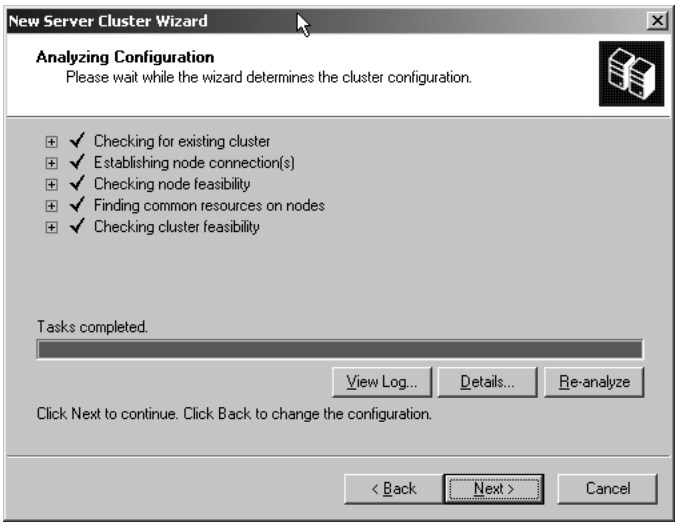

Figure 6.11 Cluster configuration analysis has succeeded.

The next dialog box prompts you for an IP address that Cluster Administrator can connect to. This is shown in Figure 6.12. Enter the IP address and click Next.

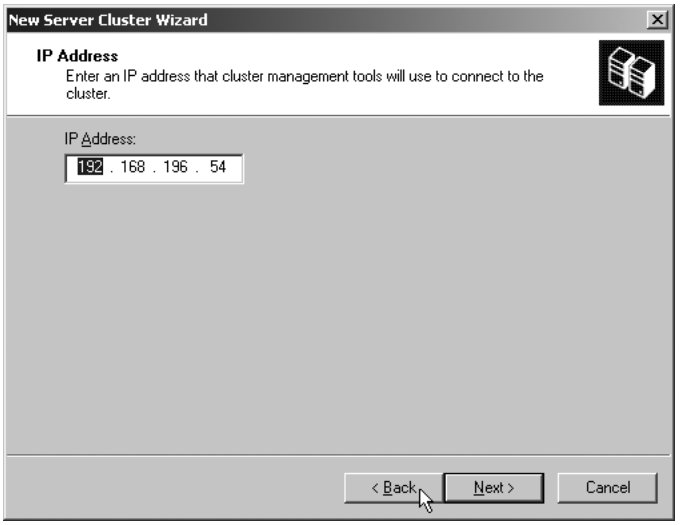

Figure 6.12 Cluster configuration IP address requirement.

The Cluster Service Account dialog box now appears. This is shown in Figure 6.12. You need to enter an account name, its password, and domain before continuing. Create an account specially for the cluster services account (create a separate account for each cluster). See Figure 6.13 for creating a cluster service account.

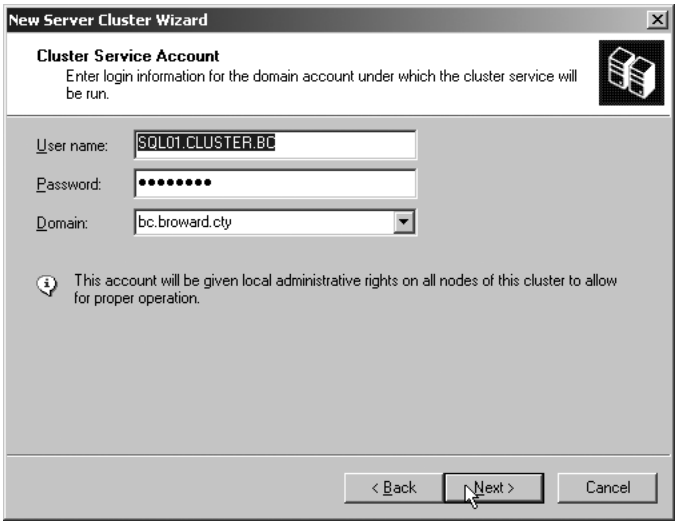

Figure 6.13 Cluster service account configuration requirement.

In the example shown, it is clear from the account name that the cluster service account is intended for the first SQL Server cluster. Under no circumstance make the account a member of Domain Admins. Making the cluster service account a member of Domain Admins was a common practice with earlier version of Windows. With Windows Server 2003, the account only needs to have administrative rights on each knot of the cluster. Upon entering the account data, click Next. The proposed cluster configuration is presented in the next dialog box. You can confirm the configuration and then go back to make a change if needed. If everything checks out, then click Next to begin the installation. Upon successful creation of the cluster, the dialog box shown in Figure 6.14 appears. When you again have a solid green line in the progress bar, you have yourself a cluster. The next dialog box gives you an opportunity to examine the cluster installation log.

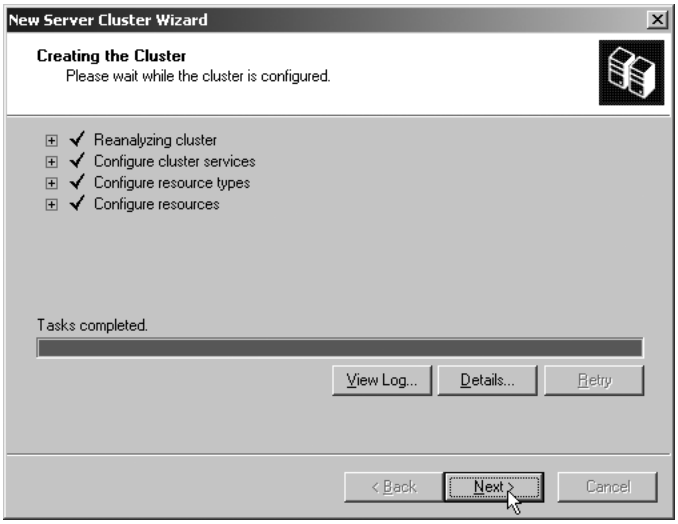

Figure 6.14 The cluster has been successfully created.

Now you can close Cluster Administrator and then reopen it to attach to the local node where you now have a single node cluster running. You attach Cluster Administrator to the name of the cluster or you can use the . (dot) notation, which is the symbol for local. If the Administrator attaches to the node successfully, the cluster can be accessed and your configuration can continue. This is demonstrated in Figure 6.15.

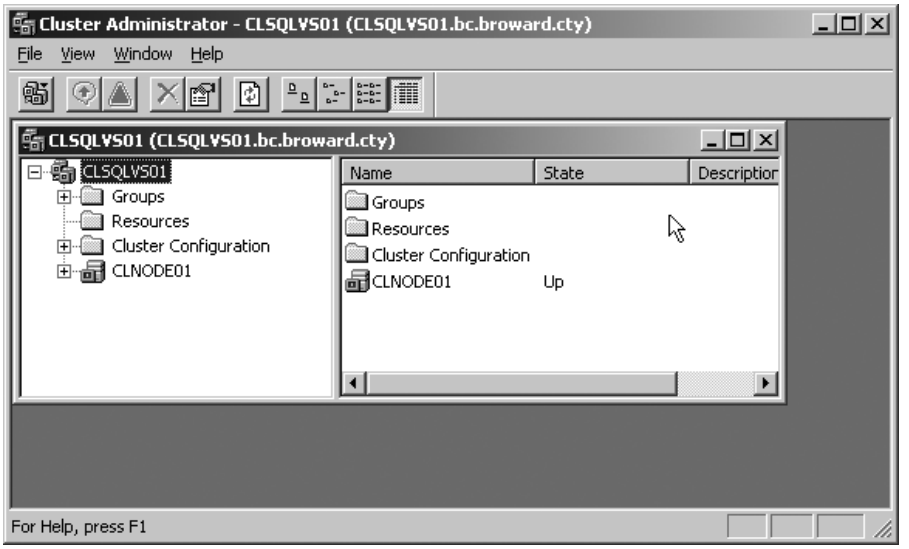

Figure 6.15 Attaching to the cluster.

You can now continue to build the cluster by adding additional nodes to it. In Cluster Administrator, click File, select New, and then select Node from the child menu. This is shown in Figure 6.16.

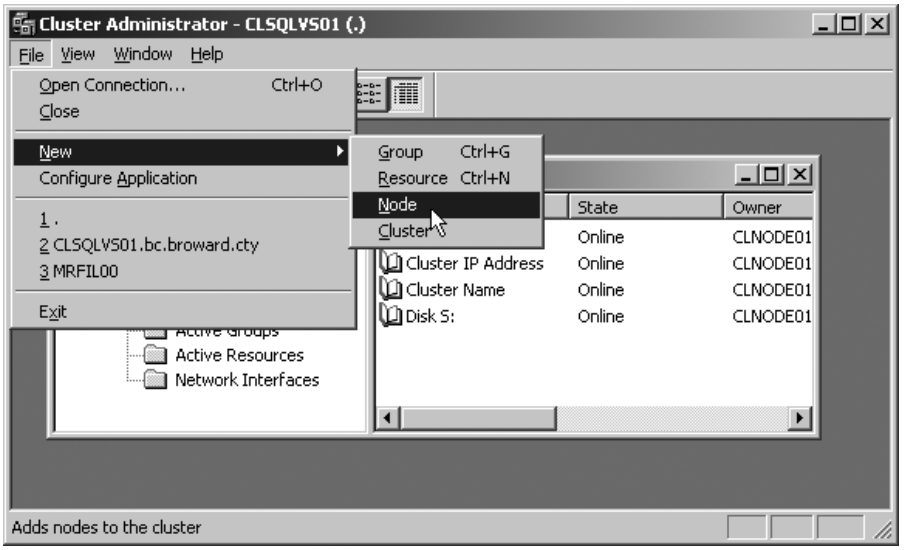

Figure 6.16 Adding nodes to the cluster.

Upon selecting New, the Add Nodes Wizard appears and prompts you to enter the name of the server that will be added as a new node to the cluster. This dialog box is shown in Figure 6.17.

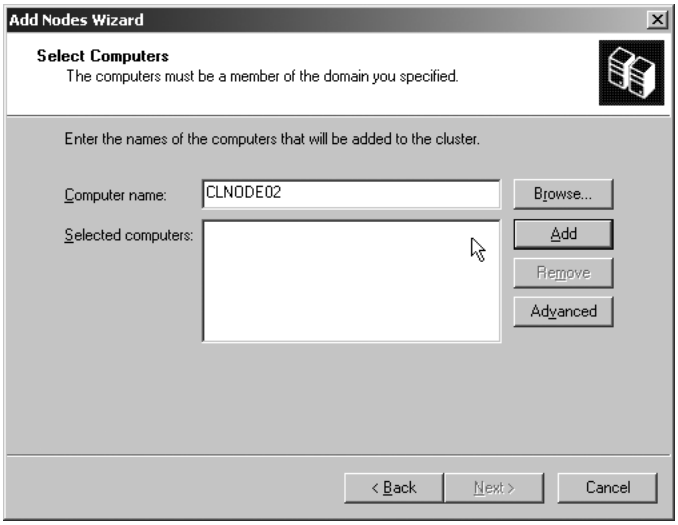

Figure 6.17 The Add Nodes Wizard.

Add the computer details and click Next. From now until the end, the process is the same as before for the first node. You are asked again for cluster service account information, and you have to provide the same service account used for the first node in the cluster. The Add Nodes Wizard again performs Configuration Analysis. When you see a green progress bar, you have a two-node cluster and you are ready to begin configuring resources for the cluster.

Cluster Administrator can tell you whether the cluster is operating properly. Open a command prompt and enter the command **cluster.resource**. This action lists the status for the available resources of the cluster. (You can issue this command even before you have added the second node to the cluster.) This is illustrated in Figure 6.18.

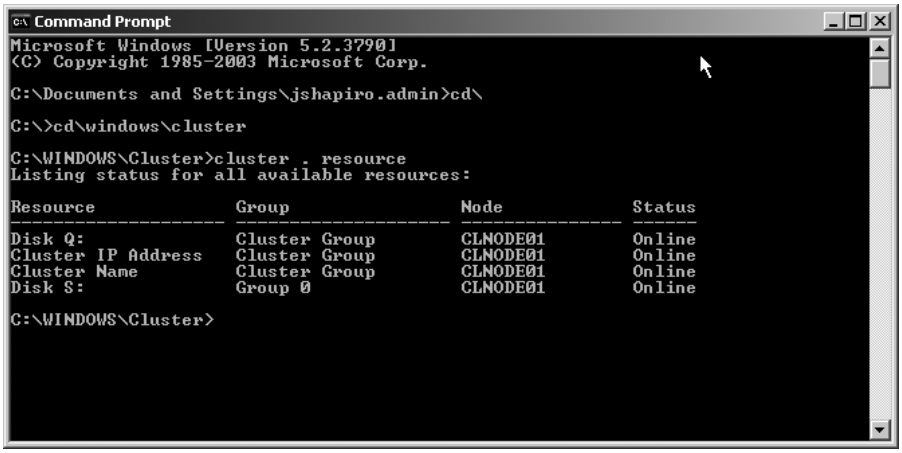

Figure 6.18 Checking cluster resource status.

It is also important to check whether the cluster has been registered in DNS and can be accessed from the network. You can do this by simply pinging the cluster name from the command line as demonstrated in Figure 6.19.

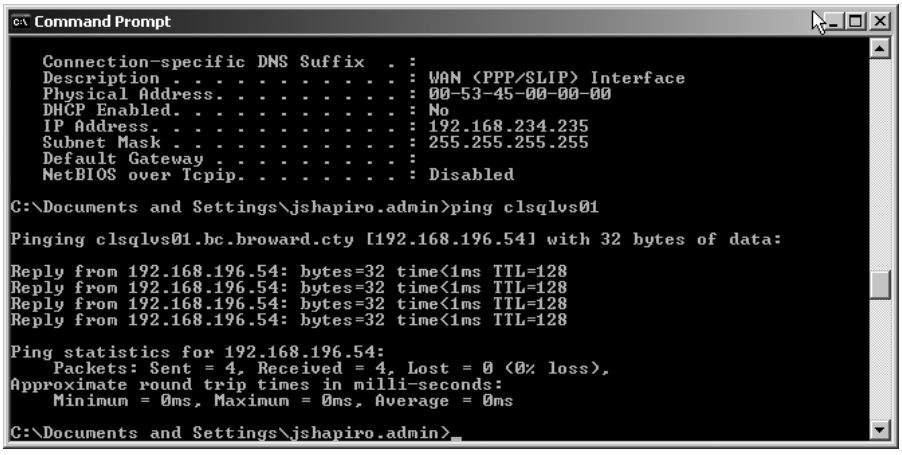

Figure 6.19 Ping the virtual server or cluster name on the network.

During the cluster creation process (using the Quorum button on the Proposed Cluster Configuration page), you are able to select a quorum resource type (that is, a Local Quorum resource, Physical Disk, or other storage class device resource, or Majority Node Set resource).

Troubleshooting

If things go bad and the cluster fails, you can simply back out of the clustering process, fix the errors, and restart the process. Usually the clustering process simply starts again with no issues. It is possible to corrupt the cluster database or contaminate it with invalid data. You may have to back out a node from the cluster, and it may not be possible to do this cleanly.

If you need to evict a node from the cluster, you can do this from the Cluster Administrator. Figure 6.20 shows the process of evicting a node that has for some reason become inoperable.

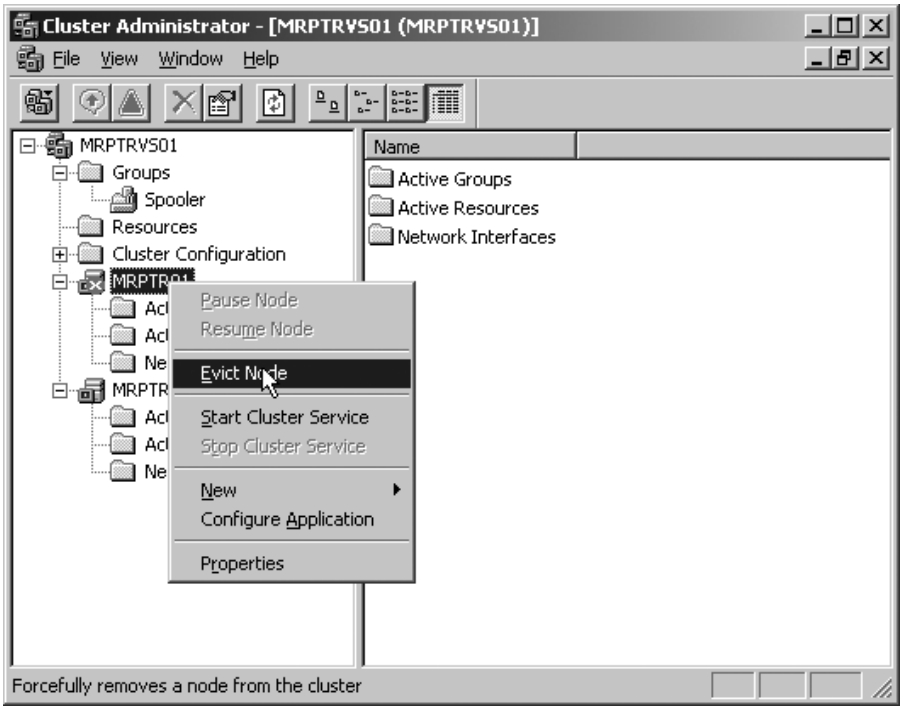

Figure 6.20 Evicting a node from the cluster.

Now, if the database is corrupt, it might not be possible to evict the node, and you may have to destroy the cluster and start all over again. When you shut down a cluster, you do not remove the cluster database (it's like the WINS or DHCP database; it's always there). It remains on the disk and it *can* remain in corrupt state. You are unable to re-create a cluster until the database is clean again. If you have cause to blow away the cluster and start all over again with a clean database, then perform the following steps.

Open up the command window on each node and change the directory to the Cluster folder in the system root (such as C:\Windows\Cluster). Then run the **/forcecleanup** command. The exact command you use is very important and not easily remembered. See Figure 6.21, which demonstrates this.

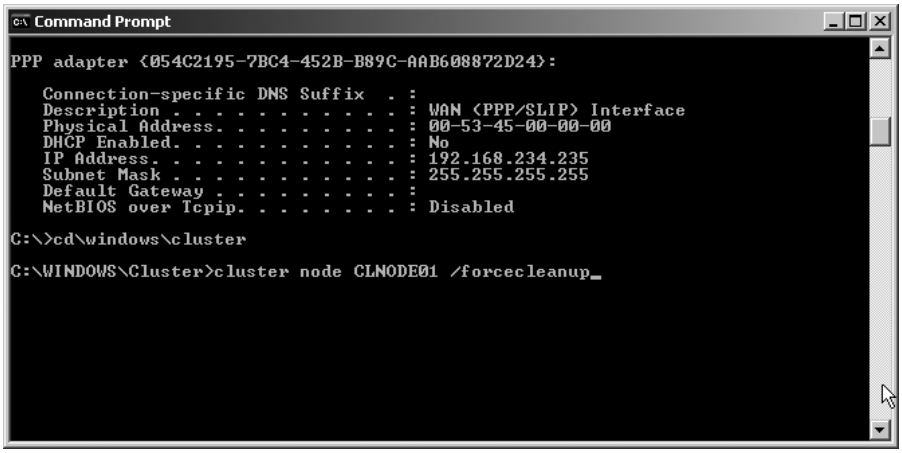

Figure 6.21 Cleaning up the cluster database on a node.

Now you have seen the Cluster command used for more than one reason. It is obviously clear you can call the Cluster executable from a script and configure a cluster after an unattended setup. As soon as the operating system is online, you can run a script to invoke the cluster **/create** command and supply it the the necessary configuration parameters at command line. Imagine that you drop a CD into a blank server and go have a cup of coffee. When you return, you have a cluster running and serving tens of thousands of users.

Time-Out

So, what was accomplished in this chapter? We did a few things. We set up a network around a single support server used for updates, patches, tools, installation bits, and DHCP services. We then set up the root domain controllers to create the forest and extended the forest to accommodate Exchange 2003. We performed quality assurance and set the scene to implement the child domain. Remember all high availability systems are installed into the child domain as are Exchange 2003 servers. After the network was established, we installed the cluster nodes and established the clusters. Let's now move onto the next chapter, which begins with configuring the applications that run on these clusters.