

Microsoft® Solution for Systems Architecture Internet Data Center

Version 1.5

Chapter 5 Building the SQL Server 2000 Clusters

Prescriptive Architecture Guide

Abstract

Setting up the Windows Clustering feature with the Microsoft® Windows® 2000 operating system and Failover Clustering with Microsoft SQL Server™ is a lengthy process involving installation of various software components in a specific order. It also involves adequate planning of hardware and network configuration. This chapter aims to assist the reader in gaining a basic understanding of the process and provides a guide to the installation and configuration of a clustered environment.

Copyright © 2002 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

Trademark Information

EMC2, EMC, and Symmetrix are registered trademarks and EMC Enterprise Storage, The Enterprise Storage Company, The EMC Effect, Connectrix, CLARiiON, EMC ControlCenter, ESN Manager, and EMC Navisphere are trademarks of EMC Corporation.

Microsoft, Windows, Windows NT, Active Directory, ActiveX, JScript, NetMeeting, SQL Server, and Visual Basic are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

CONTENTS

INTRODUCTION	1
Audience	1
Resource Requirements	1
Hardware Resource Requirements	2
Software Resource Requirements	2
System Prerequisites	2
Design Considerations	2
INSTALLATION STEPS.....	4
Determining System Requirements	4
Configuring Shared Storage	5
Pre-installation Procedure	5
Configuring the Network for Clustering	5
Configuring Network Adapters	5
Verifying Network Connectivity	7
Verifying Domain Membership	7
Setting Up a Cluster User Account	7
Setting Up Shared Disks	7
Partitioning Shared Disks	7
Verifying Disk Access	8
Configuring the First Node	8
Configuring the Second Node	10
Verifying Installation	10
Verification A: Failover of Resources	10
Verification B: Resource Access	10
Verification C: Taking a Node Off the Network	11
Verification D: Powering Off a Node	11
Installing SQL Server Failover Clustering	11
Configuring DTC	11
Creating a New SQL Server Cluster	11
Upgrading the Servers in a Cluster	14
Managing the Servers in a Cluster	14
Administering SQL Server in a Cluster Configuration	14
Upgrading SQL Server	14
Creating an Emergency Repair Disk	14
Improving System Performance	15
Best Practices	15
SUMMARY	16
Additional Information	16

INTRODUCTION

The Microsoft® Systems Architecture (MSA) Internet Data Center (IDC) design includes a high-availability Data and Management Virtual Local Area Network (VLAN) with several server configuration options, all of which enhance application availability. To create this high-availability VLAN, you can add hardware to each server in the Data and Management VLAN to handle increased application load. This approach is more cost-effective and easier to implement than other options; however, each server remains a single point of failure. Another option is to partition the data-tier component of your application across a number of servers, thereby distributing the load. For example, if you were to put customer data on one server and order data on another, you could potentially double the application speed at the data layer. The results achieved thereby could be significant, but would involve extensive application modifications.

A third option, building out a virtual server across a pair of servers that are running Windows® Clustering and SQL Server™ Failover Clustering, is easier to implement, offers significant benefits, and requires little, if any, application modification. These clustering technologies work together to provide seamless, automated restarts of the application after hardware failure. The cluster removes any single point of failure from the data tier and provides scalability. To implement this option, you install Windows Clustering and SQL Server Failover Clustering on top of the appropriate hardware. This chapter details the steps required to implement this last option.

Audience

It is recommended that all management and operations personnel involved in the planning, configuration, and maintenance of Windows Clustering and SQL Server Failover Clustering familiarize themselves with the concepts outlined in this document. It is further recommended that all personnel involved in the architecture and development of applications that rely on a clustered resource, read and understand this guide.

Resource Requirements

The following roles are required for a successful cluster implementation:

- Server administrator
- Operating system administrator
- Network administrator
- Shared disk/network storage administrator
- Windows Clustering administrator
- SQL Server administrator
- SQL Server Failover Clustering administrator
- Backup administrator

While your organization may not have a separate person fulfilling each role, it is important to account for specific considerations associated with each role during the planning and execution of the cluster. The planning and execution of the cluster must include a highly skilled engineer in each required role.

Hardware Resource Requirements

The following hardware is required for a successful cluster implementation:

- Two cluster nodes. Each node must support multiple CPUs, multiple local hard drives, at least three Network Interface Cards (NICs), and a mechanism for interfacing with the shared storage resource. The nodes must also meet all system requirements associated with the Microsoft Windows 2000 Advanced Server operating system and SQL Server 2000 and be compliant with the Hardware Compatibility List (HCL).
- A shared storage resource.

Software Resource Requirements

The following software is required for a successful cluster implementation:

- Windows 2000 Advanced Server
- SQL Server 2000 Enterprise Edition
- Emulex drivers for the Host Bus Adaptors (see Chapter 4)

System Prerequisites

The following are the prerequisites for your system before you introduce the cluster to your production environment:

- All hardware must be HCL-compliant and certified by the OEM to work with Windows Clustering and SQL Server Failover Clustering.
- You must test your application against the cluster. Your application must also be able to function adequately even in the case of a service failover.
- If you opt to use the multiple-instance failover (formerly known as Active/Active) operational mode, you need to plan for additional capacity and your application must also include a throttle (a mechanism that reduces traffic to the cluster).

Design Considerations

A two-node cluster built on three network connections and configured to run in single-instance failover mode efficiently provides a high level of availability and resiliency. This configuration satisfies most scenario requirements.

The following considerations lead to the two-node cluster design:

- **Two nodes are cost-effective.** It is best to minimize the amount of hardware that sits idle while the cluster is in operation. The benefit of extra nodes decreases as the number of servers increases. In a two-node cluster, a single server sits idle, handling only private cluster traffic. In a four-node cluster, three servers sit idle. It is much easier to justify the cost of a two-node cluster, because it provides nearly the same failover capacity per service as a larger cluster.
- **Two network links promote high availability and ease of management.** With two network connections, it is possible to dedicate a network to incoming traffic (client requests), and a network to cluster traffic. Losing one of these connections would force incoming and administrative traffic to ride on the same, highly burdened, and possibly unsafe link.
- **Single-instance failover mode makes it easy to integrate the cluster with your application.** The application just needs to handle or tolerate the service restart that occurs during cluster changeover. However, if you prefer multiple-instance failover operational mode, application load concerns must be addressed prior to the installation. It is likely that additional node capacity will be required or that you will need to introduce a throttling mechanism into the application environment.

INSTALLATION STEPS

The Windows Clustering installation procedure in this chapter is based on the *Step-by-Step Guide to Installing Windows Clustering*, which is available at:

<http://www.microsoft.com/windows2000/techinfo/planning/server/clusterrsteps.asp>

Determining System Requirements

The Internet Data Center architecture uses two quad-processor Dell PowerEdge 8450 for the back-end database cluster. This hardware requirement is based on the performance-testing results of various hardware configurations.

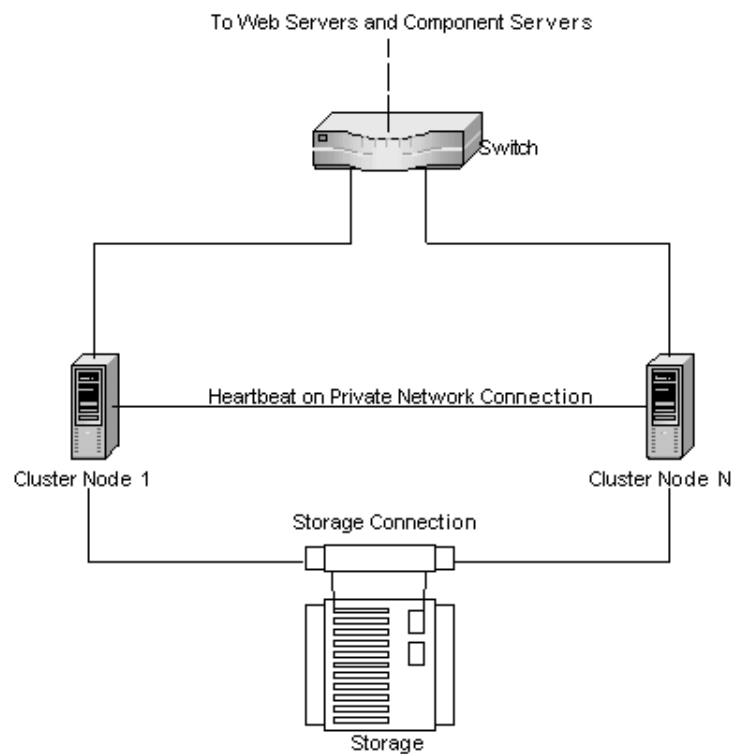


Figure 5.1 Hardware configuration for the back-end database cluster

Each database server is installed with internal storage for the operating system and access to external storage for files used by Windows Clustering and the application.

Each server has two network interface cards (NICs).

One is used for private cluster communication between the two servers to provide the cluster heartbeat. This communication can happen over an Ethernet crossover cable, or through a hub or switch. Using a switch provides the following advantages:

-
- If a node fails or is brought offline, the other node will not receive a carrier error on the heartbeat network, which would generate alerts.
 - If you are using Microsoft Windows 2000 Datacenter Server, you can add nodes to the cluster to increase capacity or consolidate services.

Another NIC is connected to the back-end network to provide access to the rest of the infrastructure. This NIC is configured for failover teaming to allow for higher availability. This NIC also acts as a redundant backup connection for heartbeat-cluster communication.

Configuring Shared Storage

- The Internet Data Center architecture employs an EMC Clarion storage server for shared storage. This server achieves high availability by clustering discrete nodes within a common enclosure and a common management environment. You can configure an EMC Clarion storage server to provide file storage for any number of servers. See Chapter 3 (Storage Area Network) for the specific configuration steps.

Pre-installation Procedure

Before you begin to install Windows Clustering, read "Installation Overview" in the *Step-by-Step Guide to Installing Windows Clustering*, which is available from the following Web site:

<http://www.microsoft.com/TechNet/win2000/clustep.asp>

Configuring the Network for Clustering

Before you begin, turn off the shared storage device and restart both database servers for network configuration.

Note Each database server must be equipped with two Network Interface Cards (NICs).

Configuring Network Adapters

Here are the basic steps for configuring Network Adapters. (For more information, see the sections entitled "Configuring the Private Network Adapter," "Configuring the Public Network Adapter," and "Rename the Local Area Network Icons" in the *Step-by-Step Guide to Installing Windows Clustering*.)

To rename the local area network (LAN) icons:

1. Determine which network segment is connected to which network connection by tracing or unplugging the cables.
2. Rename the network connections according to the corresponding network segments.

To set up TCP/IP properties for a private cluster connection:

1. On the **Wins** tab, click **Disable NetBIOS over TCP/IP** for the private cluster connection, and then click **OK**.
2. Repeat this step for the second database server's set of Transmission Control Protocol/Internet Protocol (TCP/IP) properties.

Here is an example of TCP/IP properties for the two database servers.

	Database Server 1	Database Server 2
Back-End Connection <ul style="list-style-type: none">• IP address (teamed)• Subnet mask• Connection type of NIC	192.168.12.201 255.255.255.0 100BaseTx Full Duplex	192.168.12.202 255.255.255.0 100BaseTx Full Duplex
Heartbeat connection <ul style="list-style-type: none">• IP address• Subnet mask• NetBIOS over TCP/IP	192.168.14.2 255.255.255.0 Disabled	192.168.14.4 255.255.255.0 Disabled
Cluster Management <ul style="list-style-type: none">• IP address• Subnet mask	192.168.12.211 255.255.255.0	192.168.12.211 255.255.255.0
SQL Virtual Server <ul style="list-style-type: none">• IP address• Subnet mask	192.168.12.212 255.255.255.0	192.168.12.212 255.255.255.0

Verifying Network Connectivity

Verify all network connections as follows:

- Use the **IPCONFIG** command at the command prompt to view a list of network adapters with their corresponding TCP/IP settings.
- Use the **PING** command to validate all network connections.

For more information about this procedure see "Verifying Connectivity and Name Resolution" in the *Step-by-Step Guide to Installing Windows Clustering*.

Verifying Domain Membership

Ensure that all nodes in the cluster are members of the same domain and can access the domain controller. To set up domain membership for the database server, join the Internet Data Center domain from the **Network Identification** tab of **My Computer/Properties**.

For more information about this procedure, see "Verifying Domain Membership" in the *Step-by-Step Guide to Installing Windows Clustering*.

Setting Up a Cluster User Account

You must create a domain user account under which you can run Windows Clustering *before* starting the installation, because setup requires a user name and password. For more information about this procedure, see "Setting Up a Cluster User Account " in the *Step-by-Step Guide to Installing Windows Clustering*.

Setting Up Shared Disks

Ensure that all the drives have the same drive letter assigned on both nodes.

Caution

Before setting up shared disks, turn off all nodes, and then turn on only the first node and the shared storage device. If both nodes and the shared device are switched on before Windows Clustering is completely installed on the first node, you may lose or corrupt data on the shared disks.

Partitioning Shared Disks

If you have not already done so, you must create the appropriate partitioning scheme for your shared disk. Using the Disk Management feature in Microsoft Windows 2000, create the logical partitions that will reside on the shared disks. For more information about this procedure, see "About the Quorum Disk," "Configuring Shared Disks," and "Assigning Drive Letters" in the *Step-by-Step Guide to Installing Windows Clustering*.

Here are some guidelines for setting up the shared disks:

- Use a separate disk or partition, called a quorum disk, to store cluster configuration and log files that help manage the cluster.
- Format all shared disks as the NTFS file system and designate them as basic.
- Assign a drive letter to each partition.

Verifying Disk Access

It is important to test the newly created shared partitions by writing and reading a simple text document into each partition. For more information about this procedure, see "Verifying Disk Access and Functionality" in the *Step-by-Step Guide to Installing Windows Clustering*.

1. With the shared device and first node *on*, turn off the second node. You will be able to access the disk partitions from the first node
2. With the shared device *on*, turn off the first node and turn on the second one. You will be able to access the disk partitions from the second node just as you did from the first one.
3. Turn the first one back on before proceeding to configure the first node for Windows Clustering.

Configuring the First Node

Perform the following steps to configure the first node. For more information about this procedure, see the "Configuring the First Node" section of the *Step-by-Step Guide to Installing Windows Clustering*.

1. Open **Add/Remove Windows Components** from the **Control Panel**.
2. Select **Cluster Service**, and click **Next**.
3. To accept default values, click **Next**.
4. When prompted, insert the Windows 2000 CD-ROM and click **OK**. If you are installing from the network share, browse to the **Windows Advanced Server** share, click **Open**, and then click **OK**.
5. On the **Welcome to the Cluster Server Configuration Wizard** page, click **Next**.
6. On the **Hardware Configuration Information** page, click **I Understand**.
7. Click **Next**, and then click **Next** again.
8. On the **Cluster Name** page, specify a common name for the cluster, and click **Next**.
9. On the **Select an Account** page, enter the cluster user account that you set up earlier, and click **Next**.

-
10. If the cluster user account is not set up as a member of the administrators group, a dialog box appears asking if you want to make this account a member of the administrators group. Click **Yes** to add cluster user account information in the administrators group.

Note The *Step-by-Step Guide to Installing Windows Clustering* (available from <http://www.microsoft.com/windows2000/techinfo/planning/server/clustersteps.asp>) suggests making the cluster administrator log on a member of the domain administrators group (in addition to adding the log on to the local administrators group). For security reasons, adding the cluster administrator log on to this group is not recommended.

11. On the **Add or Remove Managed Disks** page, all SCSI disks not residing on the same bus as the system disk appear in the Managed Disks list. Click **Next**.
12. Select the quorum disk from the **Disks** list, and then click **Next**.
13. When the **Configure Cluster Network** page appears, click **Next** to continue.

Note The Network Connection page order may vary, depending on the network card detected by the system. However, the content remains the same.

14. On the **Network Connections** page, specify the private or heartbeat connection, select the **Internal cluster communications only (private network)** role option, and then click **Next**.
15. For the back-end connection, specify the **All communications (mixed network)** role option. This connection is used for cluster management traffic and can also serve as a redundant backup connection for the internal private cluster connection. After specifying the details, click **Next**.
16. On the **Internal Cluster Communication** page, ensure that the private cluster connection is the highest priority for internal cluster communication as it is the primary connection for heartbeat traffic. The back-end connection on VLAN 12 is of a lower priority and will only transfer heartbeat traffic if the private cluster connection fails.
17. Enter the **Cluster Management** IP address.
18. The Cluster Service Configuration Wizard will now finish setting up the clustering service with the parameters you have entered. As a last step, it will start the cluster service.

Configuring the Second Node

Ensure that the first node and shared disk device are running when you configure the second node for cluster installation. Installing Windows Clustering on the second node is much simpler than on the first node. You follow the same procedures, except for these major differences.

1. On the **Create or Join a Cluster** page, select **The second or next node in the cluster**.
2. On the **Cluster Name** page, enter the cluster name as defined for the first node. Do not select the **Connect to cluster as** check box.

Note: The Cluster Administrator window will display names of the two database server nodes (MSANODE1 and MSANODE2) on the left listing panel. At this moment, all clustered resources, such as shared disks, cluster IP address, and cluster name, are under the ownership of the first node (MSANODE1). In other words, you can now access files of the shared disks through the first node, but not the second.

This completes two-node server cluster setup with single-instance failover mode where one node has full control of all shared disk resources and the other node serves as a hot standby (ready to assume control of the shared resources automatically if the first node fails).

Verifying Installation

There are four ways to verify your cluster installation.

Verification A: Failover of Resources

One simple way to verify a successful installation of Windows Clustering is to move the cluster group of resources from one node to the other:

1. Expand the SQL cluster node.
2. Expand **Groups**.
3. Right-click **Cluster Group** in the Cluster Administrator window.
4. Select **Move Group**. Within a few seconds, you will observe that all shared resources, such as the shared disks, are moved under the ownership of the other node.
5. After performing your verification, repeat the last step to make sure SQL01 has ownership of the resources.

Verification B: Resource Access

1. Map a network drive from one of the front-end servers. For example, from one of the Web servers in the production farm, map a network drive to \\SQLCLUSTER\\S\$.
2. Create a simple Notepad document in that network drive, and make sure you can access and read it.

-
3. Fail over the server cluster manually by moving the cluster resources from one node to the other in Cluster Administrator.
 4. When all resources have been switched to the other node, see if you can access and read the Notepad document from the Web server again.

Verification C: Taking a Node Off the Network

A third (and more drastic) way to test server cluster failover is to unplug the controlling node's network cables. Before trying this one, use the first two tests to make sure the surviving node is in control of the shared resources.

Verification D: Powering Off a Node

The final (and most drastic) test is to shut off the power of the controlling node and watch for the surviving node to take over the operations automatically. Perform this test after performing *all* these tests for the Internet Data Center architecture setup, to confirm that the server cluster is fully operational.

Installing SQL Server Failover Clustering

This section explains how to install SQL Server 2000 with Failover Clustering enabled.

Configuring DTC

The server cluster will be Distributed Transaction Coordinator (DTC) cluster-aware. For more information, see the "Failover Clustering Dependencies" SQL Server Books Online document.

To configure DTC:

1. Start **COMCLUST.EXE** from the command prompt.
2. In the Cluster Administrator window, verify that MSDTC has been added to the cluster group as one of the shared resources.
3. Make sure the MSDTC resource is **Online**.
4. Repeat the same procedures in the other server node.
5. Verify failover of the DTC by moving its group twice. Make sure SQL01 owns the resource after the verification.

Creating a New SQL Server Cluster

You can use the SQL Server Installation Wizard to set up SQL Server on the local hard drives of both server nodes as well as in the shared disks. For more information, refer to the SQL Server Books Online documentation.

Note Complete the following steps in the Active Server node that is currently controlling the shared disk.

To create a new SQL Server cluster:

1. You need to create a domain user that will be used as the identity of the SQL Server services. In this example, we have called this user SQLServerService.
2. Launch the Setup Wizard from the SQL Server 2000 Enterprise Edition compact disc on the primary node (MSADSQ1).
3. Make sure the **Virtual Server** option is selected and then assign a virtual server name for the database cluster.

Note: This is a new name (for example, SQLSERVER), which is different from the administration name of the cluster (SQLCLUSTER).

4. Click **Next**.
5. Provide the required information in the **Name** and **Company Name** text boxes and then click **Next**.
6. Accept the license agreement.
7. Enter the product key, and then click **Next**.
8. Enter the **SQL Virtual Server** IP address, and then click **Next**.
9. On the **Cluster Disk Selection** page, select the drive on which you want your data to be located (S: for example), and then click **Next**.
10. On the **Cluster Management** page, make sure the two servers of the cluster are added as configured nodes, and then click **Next** to confirm the default setting.
11. Type the administrator password wherever required for both the server nodes.
12. Uncheck the **Default** box, and type the SQL Server Instance name. For example MSADNODE1.
Note A named instance of SQL Server 2000 will use a dynamic destination port by default. This port should be changed to a fixed port prior to configuring the firewall, which in this example is TCP port 1433. The SQL Server Network Utility should be used to configure the destination port. See SQL Server Books Online for information on how to use the SQL Server Network Utility.
13. Change the destination location of the data files. Do not use the default drive you used for quorum resources.
14. Enter the name and password of the SQL Server service account you created before, and then click **Next**.

Note: This account is not a Administrator of the Internet Data Center Domain. It is a Domain User Account.

-
15. When prompted for an authentication mode, select the authentication mode required by the applications that will be using this instance of SQL Server. Only select **Mixed** if you need this mode.
 16. Select client-licensing information.
 17. The **Setup Wizard** now copies the program files to both local hard drives of the two server nodes and then copies the data files into the shared disk.
 18. Set the SQL Server 2000 destination port to 1433. Select **Start, Programs, Microsoft SQL Server, Client Network Utility**. Select **TCP/IP** in the "Enabled Protocols by order window", then select **properties**. Type 1433 in the Default port and click **OK**. Stop and restart the SQL Server 2000 Service for the change to take effect. This completes the setup.

Verifying the Cluster Creation

Now it is necessary to verify that the cluster setup has been successful by using the following steps.

To verify that the installation has been successful:

1. Start the Cluster Administrator.
2. Verify that SQL Server-related resources have been added to the cluster group. This group owns the disk resource that you specified for data.
3. Right-click the **Cluster Group** icon to initiate the manual failover.
4. Select **Move Group**. This switches the ownership of all cluster resources, including the shared disks and SQL Server, from one server node to the other. Initiate the move a few times to make sure failovers occur back and forth between the two server nodes.
5. Simulate network and hardware failure by unplugging network cables from the active server node or by shutting it off.
6. Verify that failovers occur and that the surviving node is taking full control of all the cluster resources.
7. Set up an Open Database Connectivity (ODBC) connection from a third computer to the virtual database server. For example, create a database connection with server name SQLSERV and sample database named Northwind from one of the front-end Web servers.
8. Open a table with the new ODBC connection by using MS Query or similar tools.
9. Verify the connection by the returned data records of the retrieved table.

-
10. Fail over the server cluster by moving the cluster group from one node to the other. Next, verify that the ODBC connection is still active and returning the correct data set.

Upgrading the Servers in a Cluster

You can reduce the impact of doing blocking maintenance tasks on your SQL Server nodes by performing this task on the passive nodes while the active nodes service the application users. Please refer to the SQL Server documentation for more information on how to plan and execute maintenance tasks on a cluster.

Managing the Servers in a Cluster

As part of regular maintenance and management of the servers, it is recommended that you perform system backups and recovery, and monitor and improve the server performance. For information about monitoring Windows 2000-based servers, see the *Monitoring Reliability and Availability of Windows 2000-based Server Systems* white paper available from the following Web site:

<http://www.microsoft.com/windows2000/techinfo/administration/cluster/monitorrel.asp>

Administering SQL Server in a Cluster Configuration

You administer SQL Server in a cluster by using SQL Server Enterprise Manager, just as you would a non-clustered server, although the Manager always considers a clustered SQL Server instance to be remote.

Always use the Windows Clustering administrator to start or shut down SQL Server in a clustered configuration. The software monitors all starts and shutdowns of SQL Server instances in the cluster. It does not distinguish between a server failure and a user throwing the switch and, therefore, it considers any shutdown it is not managing to be a failure. It then restarts the server on the backup node.

Upgrading SQL Server

To upgrade the SQL Server hardware or software, first use the Windows Clustering administrator to move the server from the primary node to the backup node manually. Remember that when you do this, the primary server client utilities are not available until you start them on the backup node. For more information, refer to the SQL Server 2000 documentation.

Creating an Emergency Repair Disk

An emergency repair disk can greatly simplify the restore process (which can become very complex) by providing information needed for recovery. Create one, keep copies safe but accessible, and update them whenever the system configuration changes.

Note The emergency repair disk does not contain the cluster configuration. This information is stored in a separate registry hive and can be restored from a system backup, provided you ensure that the hive is included in the system backups.

Improving System Performance

You can monitor the performance of all servers in the cluster from a single location. To improve system performance, you can increase cluster server CPU power or add more servers to the cluster. You can also upgrade to Windows 2000 Datacenter Server, which allows you to cluster up to four servers.

Best Practices

Here are some best practices for implementing clusters:

- Installation order is important. Install Windows Clustering on one node and run it before installing it on a second node. *Do not* install Windows Clustering if both nodes are online, both are connected to the shared storage, and if neither is running Windows Clustering.
- Do not try to change the computer names on either node in a cluster after you have installed Windows Clustering.
- Deselect the check box for **Auto-start service when OS starts** in the SQL Server Service Manager properties dialog box on each SQL Server computer in the cluster. SQL Server should only be managed and controlled from the cluster administrator.
- Do not use WINS static entries for cluster nodes or cluster addresses.
- Do not configure cluster resources to use unsupported network protocols or related network services such as Internetwork Package Exchange (IPX), NetBIOS Enhanced User Interface (NETBEUI), Data Link Control (DLC), Appletalk, and Services for Macintosh. Windows Clustering works only with TCP/IP.
- Do not delete the HKEY_LOCAL_MACHINE \System \Disk registry key while the cluster is running or if you are using local software fault tolerance.

SUMMARY

Windows Clustering and SQL Server Failover Clustering significantly increase reliability and availability, characteristics crucial to enterprise networks. Setting them up requires that you carefully plan hardware and software configuration, then carefully follow installation processes in sequence.

Additional Information

The SQL Server 2000 Resource Kit is an excellent guide to deploying, managing, optimizing, and troubleshooting SQL Server 2000. Detailed information on the contents of this kit and sample chapters can be obtained from the following Web site:

<http://www.microsoft.com/sql/techinfo/reskit/default.asp>

Additional SQL Server 2000 documentation and white papers are available from the Microsoft TechNet SQL Server Web site:

<http://www.microsoft.com/technet/prodtechnol/sql/default.asp>

For the latest information on Windows 2000 Server, Terminal Services, and Remote Administration, visit the Windows 2000 Server Web site:

<http://www.microsoft.com/windows2000/server/default.asp>

Also refer to the *Windows 2000 Server Resource Kit* and the *Windows 2000 Server Deployment Planning Guide*.