

# Microsoft® **Exchange** 2000 **Server**

## **Deploying Microsoft Exchange 2000 Server Clusters**

Published: July 2001

Updated: May 2002, July 2002

Applies To: Exchange 2000 Server SP3

# Copyright

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

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

©2002 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, MSDN, Outlook, Windows, and Windows NT 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.

# Table of Contents

Introduction	1
Exchange 2000 Clusters	1
Windows 2000 and Exchange 2000 Version Requirements	2
Exchange Virtual Servers	2
IP Addresses and Network Names	5
Physical Disk	5
Exres.dll	5
Quorum Disk Resource	6
Planning Exchange 2000 Clusters	7
Dedicating Computers to Exchange	7
Cluster Configurations	8
Exchange 2000 Clusters on Windows 2000 Advanced Server	8
Two-Node Cluster Topology	9
Exchange 2000 Clusters on Windows 2000 Datacenter Server	9
Exchange Virtual Server Limitations	10
Fault-Tolerant Storage	11
Considerations	11
Upgrading Exchange 2000 Cluster Nodes to Exchange 2000 SP3	11
Upgrading Exchange 5.5 Cluster Nodes to Exchange 2000 SP3	11
Separate Storage Group Log Files	11
Storage Technology	11
Storage Group Limitations	11
Multiprocessor Support Limitations	12
Memory Limitations	12
Drive Letter Limitations Using Four-Node Clusters	12
An Example Exchange 2000 with SP3 Four-Node Cluster	14
Hardware, Settings, and Scenarios	14
Four-Node Cluster Illustration	16
Set Up a Two-Node Exchange 2000 Cluster	17
Requirements	17
Software Requirements	17
Shared Disk Requirements	18
Network Configuration Requirements	18

Preinstallation Information	19
Standard Installation	20
Step 1: Prepare Active Directory for Exchange 2000	20
Step 2: Install Exchange 2000 Server on Each Node	22
Step 3: Create the Exchange Virtual Servers	24
Set Up a Four-Node Exchange 2000 Server Cluster	30
Cluster Settings	30
Understanding Exchange Virtual Server and Exchange Resource Settings	31
Exchange Virtual Server Configuration Settings	31
Exchange Resource Configuration Settings	31
Exchange Logging	32
Disabling Microsoft Exchange MTA Stack Service Monitoring	33
Specifying Exchange to Log SMTP to a Shared Disk in Your Cluster	33
Configure a Clustered Back-End Server	34
Step 1: Create the HTTP Virtual Servers in Exchange System Manager	34
Step 2: Create Virtual Directories to Match the Directories Configured on the Front-End Server	36
Step 3: Add New HTTP Virtual Server Resources to the Exchange Virtual Server	37
Removing Exchange 2000 from a Cluster	38
Removing an Exchange Virtual Server from a Cluster	38
Task 1: Move All Mailboxes and Public Folder Content to Another Exchange Virtual Server	39
Task 2: Take System Attendant Resource Offline	39
Task 3: Delete Exchange System Attendant Resource	40
Task 4: Ensure the Exchange Virtual Server Server Object is Deleted from Active Directory	40
Task 5: Delete Remaining Cluster Resources	40
Remove Exchange 2000 Installation from a Cluster	41
Reliability	42
Failover	42
Storing Exchange Data	43
SMTP Queue Directory	43
.edb and .stm Files	44

Transaction Log Files	44
Performance and Monitoring	44
Memory	45
Virtual Memory	45
/3GB Switch	46
Exchange 2000 Capacity and Topology Calculator	46
Sizing Active/Passive Clusters	46
Sizing Active/Active Clusters	46
Testing Server Capacity	47
Microsoft Exchange Load Simulator	47
Exchange Stress and Performance (ESP) Tool	48
Monitoring Performance	48
Exchange 2000 Server Cluster Failover Performance	49
Extensible Storage Engine Log Checkpoint Depth	50
Microsoft Exchange Information Store Service Connections	50
SMTP Queue Size	51
Evenly Distributing Threads across SMTP, IMAP4, and POP3	51
Tuning Exchange 2000 Server Clusters	51
SMTP Percentage of Threads and Additional Threads per Processor	51
Maximum Handle Threshold	53
Disaster Recovery	53
Identifying the Cause of a Failure	54
Backing Up Data on an Exchange 2000 Server Cluster Node	54
Recovering an Exchange 2000 Server Cluster	54
Additional Resources	55
Technical Articles and Tools	55
Microsoft Knowledge Base Articles	56

# Deploying Microsoft Exchange 2000 Server Clusters

## Introduction

Windows Clustering is a Microsoft® Windows® 2000 Server feature that administrators can use to achieve continuous availability of server resources. This technical article provides information that helps administrators:

- Understand how Microsoft® Exchange 2000 Server uses Windows Clustering
- Plan for and create Exchange 2000 clusters
- Monitor the performance of Exchange 2000 clusters
- Maintain availability of Exchange 2000 clusters

You must be familiar with Windows 2000 clustering concepts before installing Exchange 2000 clusters. Many resources—including Windows 2000 Help, *Microsoft Windows 2000 Server Resource Kit*, and Web sites such as MSDN®—offer information about Windows 2000 Server clustering concepts. Additional resources are listed at the end of this document.

## Exchange 2000 Clusters

To create Exchange 2000 clusters, you must use Windows Clustering. Windows Clustering is a feature of Windows 2000 Enterprise Server and Windows 2000 Datacenter Server. Windows 2000 Cluster Service is the essential software component that controls all aspects of Windows Clustering. When installing the cluster-aware version of Exchange, Exchange 2000 automatically installs custom files and resources. When Exchange 2000 Setup is run on a node of a Windows 2000 cluster, the cluster-aware version of Exchange is automatically installed. Exchange 2000 uses the following Windows Clustering features:

- **Shared Nothing Architecture** Windows Clustering features a shared nothing architecture that does not allow dynamic load balancing, so as a result, neither does Exchange. Thus, although all nodes in the cluster can access shared data, they cannot access it at the same time. For example, if three physical disk resources are assigned to node 1 of a two-node cluster, node 2 cannot access those three disk resources until node 1 is taken offline, fails, or the disk resources are moved to node 2 manually.
- **Resource DLL** Windows 2000 communicates with resources in a cluster using a resource dynamic-link library (DLL). Exchange 2000 provides its own custom resource DLL (named Exres.dll) to communicate with Cluster Service. Communication between Cluster Service and Exchange 2000 is customized to provide all Windows Clustering functionality.
- **Groups** Exchange 2000 uses Windows 2000 cluster groups to create Exchange Virtual Servers in a cluster. An Exchange Virtual Server in a cluster is a Windows 2000 cluster group containing cluster resources, such as an Internet Protocol (IP) address and Exchange 2000 System Attendant.

- **Resources** When you install Exchange 2000 in a cluster, it adds its own resources, such as the System Attendant resource. Exchange 2000 also uses Cluster Service resources, such as the IP address, network name, and physical disk resource.

## Windows 2000 and Exchange 2000 Version Requirements

Specific versions of Windows 2000 and Exchange 2000 are required to create Exchange 2000 clusters. Table 1 outlines these requirements.

**Table 1** Windows 2000 and Exchange 2000 version requirements

Windows 2000	Exchange 2000	Exchange clusters available
Windows 2000 Server or Microsoft Windows® 2000 Advanced Server	Microsoft Exchange 2000 Server	None
Windows 2000 Server	Microsoft Exchange 2000 Enterprise Server with Service Pack 3 (SP3)	None
Windows 2000 Advanced Server	Exchange 2000 Enterprise Server	Up to two-node
Windows 2000 Advanced Server	Exchange 2000 Enterprise Server with Service Pack 3 (SP3)	Up to two-node
Microsoft Windows 2000 Datacenter Server	Exchange 2000 Enterprise Server with Service Pack 3 (SP3)	Up to four-node

**Note** Exchange 2000 Enterprise Server is required to create Exchange 2000 clusters. Windows 2000 Advanced Server supports two-node clusters, and Windows 2000 Datacenter Server supports two-node, three-node, and four-node clusters. Exchange 2000 Enterprise Server, with at least Service Pack 1 (SP1) installed, is required to create clusters on Windows 2000 Datacenter Server. You should install the latest service pack for both Exchange 2000 and Windows 2000. Consult the service pack documentation before installing Exchange 2000 and Windows 2000 service packs.

## Exchange Virtual Servers

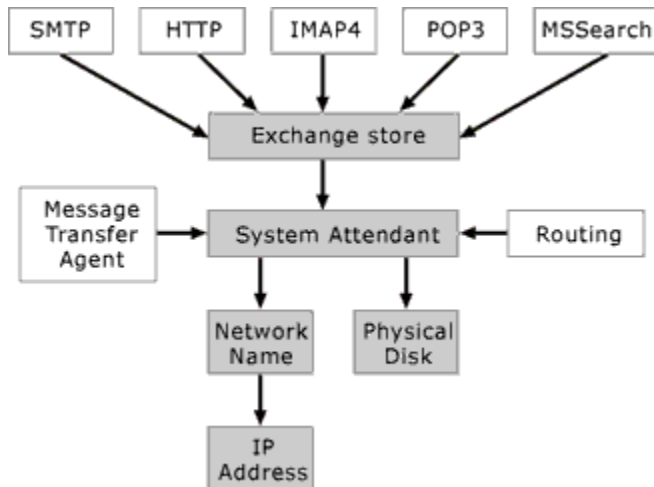
To create an Exchange 2000 cluster, you create a Windows 2000 cluster group and then add specific resources to it. The logical servers created by Exchange 2000 clusters are referred to as Exchange Virtual Servers. Unlike a stand-alone (non-clustered) computer running Exchange 2000, an Exchange Virtual Server is a cluster group that can be failed over if the server itself fails. When one computer in the cluster fails, one of the remaining nodes in the cluster takes over for the failed Exchange Virtual Server and clients can access this server using the same Exchange server name.

An Exchange Virtual Server is a cluster group that requires, at a minimum, the following resources:

- Static IP address

- Network name
- One or more physical disks on the shared storage
- An Exchange 2000 Server System Attendant resource (the System Attendant resource will install other required Exchange resources)

Figure 1 illustrates the resources of an Exchange 2000 cluster and the resource dependencies.



**Figure 1** Exchange 2000 resources and dependencies

Client computers connect to an Exchange Virtual Server the same way they connect to a stand-alone computer running Exchange 2000. Windows 2000 provides the IP address resource, the Network Name resource, and disk resources associated with the Exchange Virtual Server; Exchange 2000 provides the System Attendant resource and other required resources. When you create the System Attendant resource, all other required and dependant resources are installed.

Table 2 illustrates the Exchange 2000 components and their cluster functionality.

**Table 2** Exchange 2000 Server components and their cluster functionality

Component	Functionality	Notes
Simple Mail Transfer Protocol (SMTP)	Active/Active	Provides connections to client computers and is dependent on the Exchange store.
Internet Message Access Protocol (IMAP4)	Active/Active	Provides connections to client computers and is dependent on the Exchange store.
Post Office Protocol version 3 (POP3)	Active/Active	Provides connections to client computers and is dependent on the Exchange store.

Hypertext Transfer Protocol (HTTP)	Active/Active	Provides connections to client computers and is dependent on the Exchange store.
Exchange MS Search Instance	Active/Active	The MS Search resource provides content indexing for the Exchange Virtual Server and is dependent on the Exchange store.
Exchange store	Active/Active	Provides mailbox and public folder storage for Exchange 2000 Server. The Exchange store is dependent on System Attendant.
Message transfer agent (MTA)	Active/Passive	The MTA resource is active/passive; there can be only one MTA per cluster. The MTA is created on the first Exchange Virtual Server. All additional Exchange Virtual Servers are dependant on this MTA. The MTA serves all Exchange Virtual Servers in the cluster as long as it is online. The MTA is dependent on System Attendant.
Routing service	Active/Active	Builds the link state tables and is dependent on System Attendant.
System Attendant	Active/Active	System Attendant is the fundamental resource that controls the creation and deletion of all the resources in the Exchange Virtual Server. System Attendant is dependant on the Network Name and physical disk resources.

Exchange Server 2000 clusters do not support the following Exchange 2000 components:

- Microsoft Active Directory® Connector (ADC)
- Chat Service
- Microsoft® Exchange 2000 Conferencing Server
- Instant Messaging
- Key Management Service
- Exchange Calendar Connector
- Exchange Connector for Lotus cc:Mail
- Exchange Connector for Lotus Notes
- Exchange Connector for Microsoft Mail
- Exchange Connector for Novell GroupWise
- Microsoft Exchange Event service
- Site Replication Service (SRS)
- Network News Transfer Protocol (NNTP)

**Note** Although Exchange 2000 clusters do not support NNTP, the NNTP service (a subcomponent of the Windows 2000 Internet Information Services (IIS) component), is still a required prerequisite for installing Exchange 2000 in a cluster. After installing Exchange 2000 in a cluster, the NNTP service is not functional.

### IP Addresses and Network Names

A typical installation of a two-node cluster includes a public network used by clients to connect to Exchange Virtual Servers and a private network for cluster node communication. A typical two-node cluster has, at a minimum, seven IP addresses and five NetBIOS names and assumes the following configuration:

- Each node of the cluster has two static IP addresses (the public and private network connection IP addresses of each physical member server) and one NetBIOS name.
- The cluster itself has a static IP address and a NetBIOS name.
- Each Exchange Virtual Server has a static IP address and a NetBIOS name.

**Important** It is strongly recommended that you use a private cluster network and static IP addresses in any Exchange 2000 cluster deployment.

It is possible to deploy clustering using only a public network or Dynamic Host Configuration Protocol (DHCP) to assign and renew cluster node IP addresses, but this is not recommended. If your DHCP server is unable to renew the public IP addresses (for example, if your DHCP server fails), clients will not be able to connect to the cluster and the entire cluster may fail. Also, if your public network fails, your nodes cannot communicate, resulting in your cluster resources inability to fail over to another node.

### Physical Disk

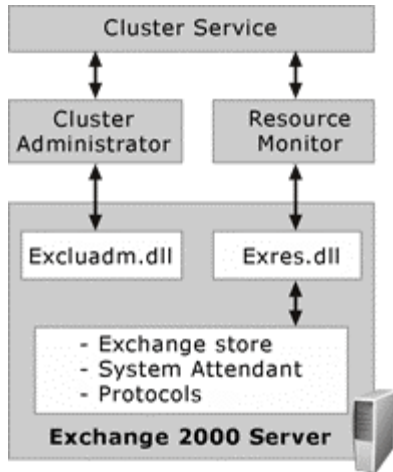
The physical disk resource is the shared disk in the cluster. You can find more information about shared disk requirements in the "Set Up a Two-Node Exchange 2000 Server Cluster" section later in this document.

### Exres.dll

Exres.dll is the Exchange 2000 resource DLL. Cluster Service communicates through a resource monitor to Exres.dll, which then communicates with the proper Exchange components. Exres.dll brings resources online and offline, checks resources with **IsAlive** calls, reports failures, and performs other component communications.

**Note** Excluadm.dll provides Cluster Administrator with the Exchange-specific administrative user interface.

Figure 2 illustrates the communication between Cluster Service and the Exres.dll resource DLL.



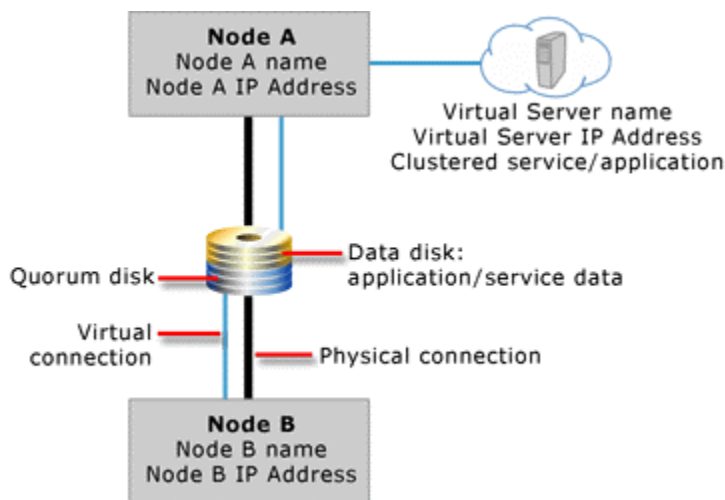
**Figure 2** Communication between **Cluster Service** and **Exres.dll**

For more information about resource DLLs, see Windows 2000 Help. For more information about the **IsAlive** function, see the MSDN Web site at <http://msdn.microsoft.com>.

### Quorum Disk Resource

The most important disk in the cluster is the disk designated as the quorum disk resource. The quorum disk resource maintains configuration data in the quorum log, cluster database checkpoint, and resource checkpoints. The quorum disk resource also provides persistent physical storage across system failures. Because the cluster configuration is kept on this disk, all nodes in the cluster must be able to communicate with the node that owns it.

Figure 3 illustrates a quorum disk resource in a two-node cluster.



**Figure 3** The quorum disk resource

When a cluster is created or when network communication between nodes in a cluster fails, the quorum disk resource prevents the nodes from forming multiple clusters. To form a cluster, a node must arbitrate for and gain ownership of the quorum disk

resource. For example, if a node cannot detect a cluster during the discovery process, the node attempts to form its own cluster by taking control of the quorum disk resource. However, if the node does not succeed in taking control of the quorum disk resource, it cannot form a cluster.

The quorum disk resource stores the most current version of the cluster configuration database in the form of recovery logs and registry checkpoint files. These files contain node-independent storage of cluster configuration and state data. When a node joins or forms a cluster, Cluster Service updates the node's private copy of the configuration database. When a node joins an existing cluster, Cluster Service retrieves the configuration data from the other active nodes.

Cluster Service uses the quorum disk resource recovery logs to:

- Guarantee that only one set of active, communicating nodes is allowed to operate as a cluster
- Enable a node to form a cluster only if it can gain control of the quorum disk resource
- Allow a node to join or remain in an existing cluster only if it can communicate with the node that controls the quorum resource

**Note** You should create new cluster groups for Exchange Virtual Servers and no Exchange Virtual Server should own the quorum disk resource.

## Planning Exchange 2000 Clusters

This section discusses issues you should consider when planning your Exchange 2000 clusters. Before you create Exchange 2000 clusters, you must create the Windows 2000 clusters. For information about how to install Cluster Service, see the technical paper *Step-by-Step Guide to Installing Cluster Service*, at <http://go.microsoft.com/fwlink/?LinkId=266>.

*Microsoft Windows 2000 Server Resource Kit* provides information about planning for Windows 2000 clusters. For detailed planning information, see the "Process for Planning Your Server Clusters" section in the "Ensuring the Availability of Applications and Services" chapter of *Microsoft Windows 2000 Server Resource Kit*.

## Dedicating Computers to Exchange

In addition to Exchange 2000, your server clusters can run other applications. However, overall performance of your server is affected by running multiple applications on your server clusters and, as a result, so is the performance of Exchange. If the purpose of your clusters is to provide Exchange services to your users, it is recommended that you run only Exchange 2000 on your clusters, and run other applications on separate hardware.

The cluster nodes of an Exchange 2000 cluster must be member servers of a domain. Exchange 2000 clusters do not support cluster nodes as domain controllers or global catalog servers.

For more information about Exchange 2000 performance, see the "Performance and Monitoring" and "Additional Resources" sections later in this document.

## Cluster Configurations

The following sections discuss Exchange 2000 cluster configurations for Windows 2000 Advanced Server and Windows 2000 Datacenter Server. Before you create your Exchange 2000 clusters, determine the level of availability expected for your users. After you make this determination, configure your hardware to the Exchange 2000 cluster that best meets your needs.

### Exchange 2000 Clusters on Windows 2000 Advanced Server

You can install either two-node active/passive or two-node active/active Exchange 2000 clusters on Windows 2000 Advanced Server.

- **Two-Node Active/Passive**

In active/passive clustering, the cluster includes one primary node and one secondary node. In an active/passive configuration, the primary node of the Exchange 2000 cluster supports all client computers while its secondary node is a dedicated server that is ready to be used whenever a failover occurs on the primary node.

If the primary node fails, the secondary node picks up all operations and continues to service client computers at a rate of performance that is close or equal to that of the primary node. Deploying an active/passive clustering configuration provides performance benefits. Because a secondary server is available to take over for the primary server, the performance of your Exchange 2000 is minimally affected by a failover.

If you deploy a two-node active/passive Exchange 2000 cluster, you must have no more than one Exchange Virtual Server in the cluster. It is recommended to deploy active/passive clusters rather than active/active clusters. Active/passive clusters are preferred because one node is dedicated to taking over a failed Exchange Virtual Server should the other node fail.

To create a two-node active/passive Exchange 2000 cluster, see the "Set Up a Two-Node Exchange 2000 Server Cluster" procedure later in this document.

- **Two-Node Active/Active**

Similar to active/passive clustering, in active/active clustering, when one node fails or is taken offline, the other node in the cluster takes over for the failed node. However, because the failover causes the remaining nodes to take on additional processing operations, the overall performance of your Exchange 2000 cluster may be reduced. For performance, availability, and scalability reasons, active/passive cluster configurations are recommended over active/active configurations. For more information about recommended configurations, see *Microsoft Exchange 2000 Server Resource Kit*.

In a typical active/active configuration, each node owns at least one Exchange Virtual Server. However, if the number of Exchange Virtual Servers in the cluster is greater than or equal to the number of nodes in the cluster, and all the Exchange Virtual Servers are owned by only one node in the cluster, then the cluster is still considered an active/active cluster.

**Note** You must plan for storage group limitations when creating active/active Exchange 2000 Server clusters. For more information, see "Storage Group Limitations" in the "Considerations" section later in this document.

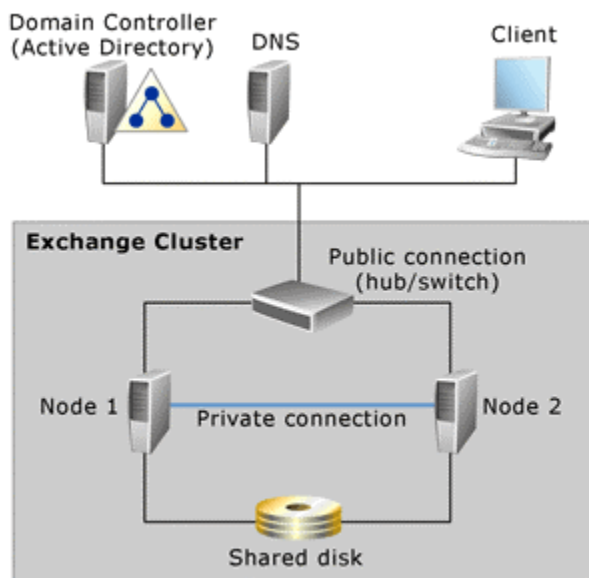
To create a two-node, active/active Exchange 2000 cluster, see the "Set Up a Two-Node Exchange 2000 Server Cluster" procedure later in this document.

**Note** To migrate from an active/active configuration to an active/passive configuration, you must remove the second Exchange Virtual Server. For information about removing an Exchange Virtual Server, see "Removing an Exchange Virtual Server from a Cluster" later in this article.

### Two-Node Cluster Topology

If the servers in your Exchange 2000 cluster are running Windows 2000 Advanced Server, you can create either an active/passive or active/active cluster. Figure 4 shows an example of a two-node cluster topology that you can use with Windows 2000 Advanced Server. Both cluster nodes are members of the same domain. The cluster nodes are connected to the public network and a private cluster network.

If only one cluster node owns one Exchange Virtual Server, this is an active/passive configuration. If both nodes own one or more Exchange Virtual Servers, or if either node owns two Exchange Virtual Servers, this is an active/active configuration (see Figure 4).



**Figure 4** Example two-node Exchange 2000 cluster

### Exchange 2000 Clusters on Windows 2000 Datacenter Server

Windows 2000 Datacenter Server increases the number of cluster nodes available on Windows 2000 Advanced Server from two to four. Exchange 2000 supports a maximum of three active nodes in an Exchange cluster configuration.

**Note** Four-node Exchange 2000 cluster configurations, in which all nodes have an active Exchange Virtual Server, are not supported on Windows 2000 Datacenter Server.

In a four-node Exchange cluster with three active nodes and one passive node, three nodes of the cluster host active Exchange Virtual Servers, while one node is available to host an Exchange Virtual Server if failover occurs. You can increase the number of passive nodes for increased reliability; for example, you can configure your Exchange cluster to have two active nodes and two passive nodes.

- **Four-Node Active/Passive with Three Active Nodes**

The four-node cluster configuration with three active nodes and one passive node includes three Exchange Virtual Servers each owned by an individual cluster node. One cluster node is passive and available to take an Exchange Virtual Server, should a failover occur.

- **Four-Node Active/Passive with Two Active Nodes**

The four-node cluster configuration with two active nodes and two passive nodes includes two Exchange Virtual Servers, each owned by an individual cluster node. Two cluster nodes are passive and available to take ownership of an Exchange Virtual Server, should a failover occur. Because this configuration includes two dedicated nodes to fail over to, this scenario may provide higher availability than the cluster configuration that includes three active nodes.

### Exchange Virtual Server Limitations

The number of Exchange Virtual Servers that can be managed by Exchange 2000 is limited. After reaching the maximum number of Exchange Virtual Servers for a specific configuration, you cannot add any additional Exchange Virtual Servers to the cluster. Table 3 lists the maximum number of Exchange Virtual Servers for each configuration.

**Important** Although it is possible to configure an active/active two-node cluster to have up to four Exchange Virtual Servers, active/active clustering is generally not recommended. It is recommended to use active/passive clustering, which allows higher availability and reliability for your clusters.

**Table 3** Maximum number of Exchange Virtual Servers per cluster

Number of nodes in a cluster	Maximum number of Exchange Virtual Servers in active/passive configurations	Maximum number of Exchange Virtual Servers in active/active configurations
2 nodes	1	4
3 nodes	2	Not supported
4 nodes	3	Not supported

**Note** If a two-node cluster with more than two Exchange Virtual Servers configured is to be upgraded to a three-node or four-node cluster, the Exchange Virtual Servers that exceed the limit must be removed.

## Fault-Tolerant Storage

The use of fault-tolerant RAID configurations is strongly recommended for all disks. If possible, use a shared RAID-1 or RAID-0+1 disk array for storage of the Exchange rich-text (.edb) and native multimedia (.stm) files and a separate RAID-1 or RAID-0+1 disk array for the transaction log files.

For more information about storage options for Exchange 2000, see "Storing Exchange Data" in the "Reliability" section later in this document.

## Considerations

The following considerations are important when creating Exchange 2000 clusters. These considerations apply to Exchange 2000 clusters on Windows 2000 Advanced Server and Windows 2000 Datacenter Server.

### Upgrading Exchange 2000 Cluster Nodes to Exchange 2000 SP3

When upgrading a node of an Exchange 2000 cluster to SP3, consider manually failing over the Exchange Virtual Server owned by that node to another node. Then upgrade the node to SP3, and manually fail back the Exchange Virtual Server. This procedure will allow users to access their Exchange Virtual Server on another node during the upgrade. It is recommended that you upgrade one node of an Exchange cluster at a time.

### Upgrading Exchange 5.5 Cluster Nodes to Exchange 2000 SP3

The procedures for upgrading your cluster nodes from Exchange 5.5 to Exchange 2000 are detailed and beyond the scope of this document. For information about upgrading to Exchange 2000, see Microsoft Knowledge Base article [Q316886](#) "HOW TO: Migrate from Exchange Server 5.5 to Exchange 2000 Server."

### Separate Storage Group Log Files

If the storage groups for an Exchange Virtual Server are configured so that the log files are on one set of physical drives and the databases on another, all of the drives must be part of the Exchange Virtual Server. That is, all of the data must be on a shared disk, and all of the physical disk resources must be part of the cluster group. This allows the log files and the storage group databases to fail over to another node if the Exchange Virtual Server goes offline.

### Storage Technology

**It is generally recommended that you use a channel-attached disk storage system for storing Exchange 2000 database files. A channel-attached disk storage system might include a small computer system interface (SCSI), Fiber Channel, or Integrated Device Electronics (IDE). Channel-attached disk storage configurations optimize performance and reliability for Exchange 2000.**

### Storage Group Limitations

Exchange 2000 is limited to four storage groups per server. This is a physical limitation and applies to each node of a cluster as well. Table 4 illustrates this limitation.

**Table 4** A two-node active/active Exchange 2000 Server cluster configuration with too many storage groups

Exchange Virtual Server	State	Storage group names
Node 1 EVS1	Active	SG1, SG2, SG3
Node 2 EVS2	Active	SG1, SG2

In Table 4, the Exchange cluster includes five storage groups. If EVS2 on node 2 fails over to node 1, node 1 will not be able to mount one of the storage groups from node 2 because node 1 will have exceeded the four storage group limitation for a single cluster node. As a result, EVS2 will not come online on node 1. If node 2 is still available, EVS2 will fail over back to node 2.

### Multiprocessor Support Limitations

Windows 2000 Datacenter Server supports 32 processors on a single server. Exchange 2000 with SP3 will only scale effectively to 8 processors on a single server. A 32-processor server should be partitioned into four, 8-processor servers using hardware partitioning. Running Exchange 2000 with SP3 on servers with more than eight processors is an inefficient use of processor resources and should be avoided.

For more information about sizing processor resources, see the "Additional Resources" section at the end of this document.

### Memory Limitations

Exchange 2000 works well with the **/3GB** switch.

**Note** Exchange 2000 requires that the **/3GB** switch be used in conjunction with Windows 2000 Advanced Server and Windows 2000 Datacenter Server on hardware with more than 1 GB of physical RAM installed. For more information, see the "/3GB Switch" and the "Additional Resources" sections later in this document.

Exchange 2000 does not support instancing (the ability to run multiple instances of an application as separate processes on the same computer) or Physical Address Extension (PAE). These restrictions limit Exchange 2000 to about 3 GB of usable memory.

Installing more than 3 GB of physical memory on a computer running only Exchange 2000 is not recommended because the additional memory does not result in a measurable improvement of server performance. However, if you are running other memory intensive applications in conjunction with Exchange 2000, such as Microsoft SQL Server™, then increasing the physical memory size beyond 3 GB can provide greater server performance.

For more information about PAE, see "PAE Design" at the Microsoft Windows Platform Development Web site at <http://go.microsoft.com/fwlink/?LinkId=8893>.

### Drive Letter Limitations Using Four-Node Clusters

Four-node Exchange 2000 clusters have an additional limitation that you must plan for before building the cluster. Windows 2000 has a 24-disk volume limitation per server. If you plan to have the majority of disks on the server as shared cluster resources, the 24-disk volume limitation applies to the entire cluster, not just to each individual node. Regardless of the number of nodes in the cluster, the maximum number of shared disks

is 23. (The reason the maximum number of shared disks is 23 and not 24 is because a disk must be reserved for the system disk on each node).

For certain configurations, you may need to disable one or more drives to make room for more shared disks in a cluster. For example, you may want to disable the CD-ROM or DVD-ROM drives on your servers. You may also want to remove the drive mapping for the Exchange Installable File System (IFS) drive, typically drive M. For information about removing the Exchange IFS drive, see Microsoft Knowledge Base article [Q305145](#) "HOW TO: Remove the IFS Mapping for Drive M in Exchange 2000 Server." Remember that maximizing the number of shared disks can reduce your ability to map drives for network share access.

**Note** Because Windows 2000 clustering does not support the use of mount points (a form of logical disk), you cannot use mount points for your Exchange 2000 shared disks. However, mount points may be used for local drives, for example, CD-ROM or DVD drives.

This disk volume limitation is a limiting factor in how an administrator designs storage group and database architecture for an Exchange 2000 cluster. Table 5 shows some sample configurations to illustrate this limitation.

**Table 5** A 3 active/1 passive architecture with three Exchange Virtual Servers, each with three storage groups

Node 1 (EVS1 active)	Node 2 (EVS2 active)	Node 3 (EVS3 active)	Node 4 (passive)
Disk 1: SMTP/MTA	Disk 8: SMTP	Disk 15: SMTP	Disk 22: Quorum
Disk 2: SG1 databases	Disk 9: SG1 databases	Disk 16: SG1 databases	
Disk 3: SG1 logs	Disk 10: SG1 logs	Disk 17: SG1 logs	
Disk 4: SG2 databases	Disk 11: SG2 databases	Disk 18: SG2 databases	
Disk 5: SG2 logs	Disk 12: SG2 logs	Disk 19: SG2 logs	
Disk 6: SG3 databases	Disk 13: SG3 databases	Disk 20: SG3 databases	
Disk 7: SG3 logs	Disk 14: SG3 logs	Disk 21: SG3 logs	

The design shown in Table 5 is very reliable: each storage group has a dedicated drive for its databases and a dedicated drive for its log files. An additional disk is used for the Exchange Virtual Server's SMTP queue directory. However, with this design, you are limited to three storage groups per Exchange Virtual Server.

The design shown in Table 6 adds an additional storage group but, to stay within the 23-disk limit, the databases of each of the four storage groups per Exchange Virtual Server are combined across two disks. The database files (.edb and .stm) of SG1 and SG2 share a common disk volume while the database files of SG3 and SG4 share a common disk volume. The benefit of this design is that you can use all four storage groups in a four-node cluster. The disadvantage of this design is that the volumes that house the

shared storage group databases may need to be very large, and if a database disk fails, two storage groups are affected instead of one.

**Table 6** A 3 active/1 passive architecture with three Exchange Virtual Servers, each with four storage groups

Node 1 (EVS1 active)	Node 2 (EVS2 active)	Node 3 (EVS3 active)	Node 4 (passive)
Disk 1: SMTP/MTA	Disk 8: SMTP	Disk 15: SMTP	Disk 22: Quorum
Disk 2: SG1 and SG2 databases	Disk 9: SG1 and SG2 databases	Disk 16: SG1 and SG2 databases	
Disk 3: SG1 logs	Disk 10: SG1 logs	Disk 17: SG1 logs	
Disk 4: SG1 logs	Disk 11: SG2 logs	Disk 18: SG2 logs	
Disk 5: SG3 and SG4 databases	Disk 12: SG3 and SG4 databases	Disk 19: SG3 and SG4 databases	
Disk 6: SG3 logs	Disk 13: SG3 logs	Disk 20: SG3 logs	
Disk 7: SG4 logs	Disk 14: SG4 logs	Disk 21: SG4 logs	

## An Example Exchange 2000 with SP3 Four-Node Cluster

A common and reliable configuration for an Exchange 2000 Server four-node cluster is a design in which three nodes are active and one node is passive. This design attempts to balance the CPU, disk, memory, and network load of the computer so that no component of the system is bottlenecked prematurely.

### Hardware, Settings, and Scenarios

The following items list the hardware, settings, and scenarios used in the four-node cluster example:

- **Computers running Exchange 2000** In this example, four computers running Exchange 2000 are configured with the following hardware:
  - Eight, 700-MHz (1 MB or 2 MB L2 cache) processors
  - 3 GB of Error Correction Coding (ECC) RAM
  - Two, 100 Mbps or 1000 Mbps network interface cards
  - RAID-1 array with two internal disks for the Windows 2000 and Exchange 2000 Server program files
  - Two, redundant 64-bit Fiber Host Bus Adapters (HBA) to connect to the Storage Area Network
- **Networks (LAN)** In this example, the following hardware is used for the network:
  - Two, 100 Mbps or 1000 Mbps network switches (full duplex)
- **Storage Area Network** In this example, a Storage Area Network is used with the following configuration:
  - Redundant fiber switch

- 104 disk spindles (Ultra Wide SCSI) with spindle speeds of 10,000 RPM or better)
- 256 MB or more read/write cache memory
- **Exchange Virtual Servers** In this example, three Exchange Virtual Servers are used, EVS1, EVS2, and EVS3.
- **Storage Groups and Databases** In this example, the storage groups and databases use the following configuration:
  - Three storage groups per Exchange Virtual Server
  - Five databases per storage group
- **Drive Letter Configuration** In this example, the drive letter configuration is based on the drive letter configuration illustrated in Table 5 of this document.
- **Storage Area Network Disk Configuration** In this example, the following Storage Area Network disk configurations per storage type (104 shared disk spindles) are used:
  - **SMTP/MTA Drives** RAID-(0+1) array made of up four spindles
  - **Storage Group Log Drives** RAID-1 array made up of two spindles
  - **Database (.edb and .stm files) Drives** RAID-(0+1) array made up of eight spindles
  - **Cluster Quorum Drive** RAID-1 array made up of two spindles
- **Exchange Virtual Server Settings (Cluster Group)** In this example, the Exchange Virtual Servers are configured with the following settings:
  - **Preferred Owners** EVS1 on Node1, EVS2 on Node2, and EVS3 on Node3
  - **Failback** Failback is disabled on all Exchange Virtual Servers so that an administrator can move the groups back at the appropriate time
- **Cluster Resource Settings** In this example, the Exchange Virtual Server resources are configured with the following settings:
  - **Possible Owners (for all resources in the group)** All nodes are possible owners.
  - **Do Not Restart/Restart (for all resources in the group)** To enable Restart, select the **Affect the group** check box with a threshold of 3 and a period of 900 seconds.
  - **Pending Timeout** All resources are set to the default 3 minutes except for the Exchange store instance, which should be increased to between 5 and 10 minutes.
- **Performance Tuning Settings** In this example, the threads shared between the SMTP service and IMAP4, and POP3 are balanced.
- **Failover Performance Configuration** In this example, two failover performance options are available:
  - **Configured for Optimum Failover Performance** Set the Extensible Storage Engine log checkpoint depth to 5 MB on all storage groups. Decrease the SMTP Max Handle Threshold from 10,000 to 1,000.

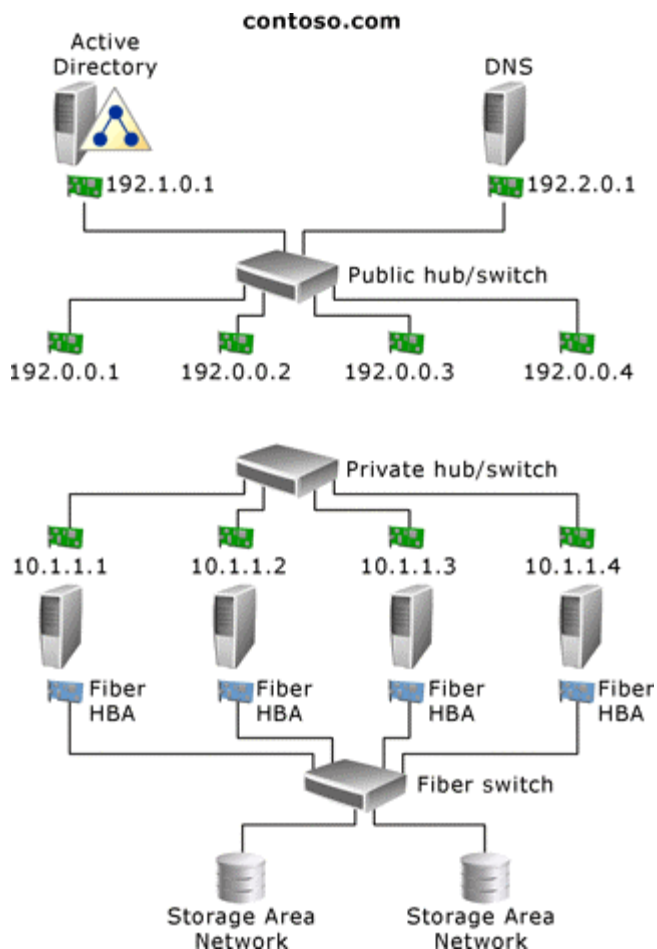
- **Configured for Optimum Run-State Performance** Leave the Extensible Storage Engine Log checkpoint depth at the default of 20 MB for all storage groups. Leave the SMTP Max Handle Threshold at the default of 10,000.

For more information, see the “Tuning Exchange 2000 Server Clusters” section later in this document.

- **Failover Scenarios** This configuration can handle a single node failure and maintain 100 percent availability after the failover has occurred. A second failure during this period leaves the cluster in a partially-up state.
- **First Failure** If node 1 fails, node 2 still owns EVS2, node 3 still owns EVS3, and node 4 takes ownership of EVS1 with all the storage groups mounted after the failover.
- **Second Failure** If another node fails, the Exchange Virtual Server will attempt to fail over to another node not hosting an Exchange Virtual Server. Because no node is available, the Exchange Virtual Server will time out and remain in a failed state.
- **Recommended Processor Load** In a four-node active/passive cluster with three active nodes, no node can own more than one Exchange Virtual Server. As a result, processor load sizing of each server can be done using Exchange 2000 Capacity and Topology Calculator, just as you would for stand-alone server deployments.

#### Four-Node Cluster Illustration

Figure 5 illustrates the four-node Exchange 2000 cluster described in this section.



**Figure 5** Example four-node Exchange 2000 cluster

## Set Up a Two-Node Exchange 2000 Cluster

The following sections contain procedures for setting up a two-node, active/passive or active/active Exchange 2000 cluster on Windows 2000 Advanced Server. These procedures assume that you have already created a two-node Windows 2000 cluster. For step-by-step instructions about creating a Windows 2000 cluster, see the technical paper *Step-by-Step Guide to Installing Cluster Service* available at <http://go.microsoft.com/fwlink/?LinkId=266>.

### Requirements

The following items are software and shared disk requirements for installing Exchange 2000 on a Windows 2000 cluster.

#### Software Requirements

The following software requirements are the minimum needed to install Exchange 2000 on a Windows 2000 cluster:

- Microsoft Windows 2000 Advanced Server (Service Pack 3 recommended)
- Microsoft Exchange 2000 Enterprise Server (Service Pack 3 recommended)

- Domain Naming System (DNS), and Windows Internet Naming System (WINS)

**Note** Terminal Services is recommended to allow remote cluster administration.

### Shared Disk Requirements

The following shared disk requirements are the minimum required to install Exchange 2000 on a Windows 2000 cluster:

- Shared disks must be physically attached to a shared bus.
- Disks must be accessible from all nodes in the cluster.
- Disks must be configured as basic, not dynamic.
- All partitions on the shared disk must be formatted NTFS.
- Only physical disks can be used as a cluster resource. All partitions on a physical disk will be treated as one resource.

### Network Configuration Requirements

Exchange 2000 is dependent on communicating with the Active Directory® directory service servers in your organization. It is important that the networks used for both client and cluster communications are configured correctly.

- **Private Cluster Network Settings** Do the following to ensure your private network is properly configured:

#### To ensure your private network is properly configured

1. In **Control Panel**, double-click **Network and Dial-up Connections**.
  2. In **Network and Dial-up Connections**, right-click **<Network Connection Name>** (where **Network Connection Name** is the name of your private network connection), and then click **Properties**.
  3. In the **<Network Connection Name> Properties** dialog box, on the **General** tab, in **Components checked are used by this connection**, ensure the **Internet Protocol (TCP/IP)** check box is selected.
  4. Select **Internet Protocol (TCP/IP)** and then click **Properties**.
  5. In the **Internet Protocol (TCP/IP) Properties** dialog box, click **Advanced**.
  6. In **Advanced TCP/IP Settings**, on the **DNS** tab, ensure that no addresses are listed in the **DNS server addresses, in order of use** list box.
  7. On the **DNS tab**, ensure that there are no suffixes listed in the **Append these DNS suffixes (in order)** list box.
  8. On the **DNS tab**, ensure the **Register this connection's address in DNS** check box is cleared.
  9. On the **WINS** tab, ensure **Disable NetBIOS over TCP/IP** is selected.
- **Public Network Settings** Do the following to ensure your public network is properly configured:

#### To ensure your public network is properly configured

1. In **Control Panel**, double-click **Network and Dial-up Connections**.
2. In **Network and Dial-up Connections**, right-click **<Network Connection Name>** (where **Network Connection Name** is the name of your public network connection), and then click **Properties**.
3. In the **<Network Connection Name> Properties** dialog box, on the **General** tab, in **Components checked are used by this connection**, ensure the **Internet Protocol (TCP/IP)** check box is selected.
4. Select **Internet Protocol (TCP/IP)**, and then click **Properties**.
5. In the **Internet Protocol (TCP/IP) Properties** dialog box, click **Advanced**.
6. In **Advanced TCP/IP Settings**, on the **DNS** tab, ensure that all of the required addresses are listed in the **DNS server addresses, in order of use** list box.
7. Ensure that the correct suffixes are listed in the **Append these DNS suffixes (in order)** list box.

### Preinstallation Information

Consider the following preinstallation information before you set up your Exchange 2000 Server cluster:

- By default, IIS is installed with Windows 2000; however, NNTP is not. NNTP is required for Exchange 2000 installation, but is not supported in the cluster. Verify that NNTP is installed before attempting to install Exchange 2000. Also verify that SMTP is installed.
- A DNS server must be available for the domain.
- All cluster nodes must be members of the same domain.
- You must have a sufficient number of static IP addresses available to you when you create the Exchange Virtual Servers. For a two-node cluster, the recommended number of static addresses is five, plus the number of Exchange Virtual Servers. For a four-node cluster, the recommended number is nine, plus the number of Exchange Virtual Servers.
- Cluster Service must be installed and running on all nodes in a cluster before installing Exchange 2000 Server. If Cluster Service is not installed and running on each node in a cluster before installation, Exchange 2000 Setup cannot install the cluster-aware version of Exchange 2000 Server.

**Note** If you installed Exchange 2000 Server before installing Cluster Service, you must uninstall Exchange 2000, install and activate Cluster Service, and then install Exchange 2000.

- Do not install Exchange 2000 on both nodes at the same time.
- The Cluster Service account must have local Administrator privileges on the cluster nodes and be a domain user account. You can establish those permissions by creating a domain user account and making this account a member of the local Administrators group on each node.
- To complete the ForestPrep phase of Exchange 2000 Setup, the installation account must be a member of the following Windows 2000 security groups: Domain Admins,

Schema Admins, and Enterprise Admins. You must run ForestPrep while you are logged on to the domain to which your schema master resides. Typically, this is the first Windows 2000 domain controller installed in your forest.

- To complete the DomainPrep phase of Exchange 2000 Setup, the installation account must be a member of the Domain Admins Windows 2000 security group. For each domain that you need to run DomainPrep, you must run DomainPrep while you are logged on to the domain to which you want to run DomainPrep.
- An Exchange 2000 Cluster server cannot be the first Exchange 2000 server to join an Exchange 5.5 site. This is because SRS is not supported on an Exchange cluster. You must first install a stand-alone (non-clustered) Exchange 2000 server into an Exchange 5.5 site prior to installing Exchange 2000 to the nodes of your cluster (the first Exchange 2000 server installed into an Exchange 5.5 site runs SRS). For information about SRS, see Exchange 2000 Help.
- For consistency, install Exchange 2000 on the same drive and in the same folder on each computer and shared storage device in the cluster. The default installation folder (for program files, which are not shared) is the local system drive.

**Note** Earlier versions of Exchange required that the Exchange program files be installed on a shared disk. This is no longer a requirement for Exchange 2000, and the program files are installed on a local drive, such as C:\Program Files\Exchsrvr, by default.

- Before you install Exchange 2000, be sure that the folder to which you will install all of the Exchange shared data on the physical disk resource is empty.
- You must install the same version of Exchange 2000 Enterprise Server on all nodes in the cluster.
- At a minimum, you must install Microsoft Exchange Messaging and Collaboration and Microsoft Exchange System Management Tools on both nodes of the cluster.

## Standard Installation

After you consider the software requirements, shared-resource requirements, and preinstallation information mentioned in the preceding sections, you are ready to set up either an active/active or active/passive Exchange 2000 cluster.

There are three steps you must perform to install Exchange 2000 in a cluster.

1. Prepare Active Directory for Exchange 2000.
2. Install Exchange 2000 on each node.
3. Create the Exchange Virtual Servers.

### Step 1: Prepare Active Directory for Exchange 2000

Perform the following tasks to prepare Active Directory for Exchange 2000.

#### **Task 1: Run ForestPrep**

Before you install Exchange 2000 anywhere in the forest, you must extend the Windows 2000 Active Directory schema. To do this task, you must run ForestPrep.

**Note** Running ForestPrep is required only if you are installing Exchange 2000 Server for the first time in your organization. If you have already installed Exchange 2000 Server in your organization, running ForestPrep is not necessary and you may skip to "Task 2: Run DomainPrep" later in this document.

1. Log on to the domain in which the schema master of your forest resides, typically the root domain. The account you use must be a member of the following Windows 2000 security groups: Domain Admins, Schema Admins, and Enterprise Admins.
2. Click **Start**, and then click **Run**.
3. In the **Open** box, type **CD\_drive\_letter\setup\i386\setup /forestprep**, where *CD\_drive\_letter* is the letter of your CD-ROM drive.
4. In the **Welcome** dialog box, click **Next**.
5. In the **End User License Agreement** (EULA) dialog box, read the EULA, and, if you agree to the terms, click **I agree**, and then click **Next**.
6. In the **Product Identification** dialog box, type the 25-digit CD key found on the back of the product CD-ROM, and then click **Next**.
7. In the **Component Selection** dialog box, verify that the action next to Microsoft Exchange 2000 is **ForestPrep**, and then click **Next**.
8. In the **Installation Type** dialog box, select **Create a new Exchange Organization**.
9. In the **Organization Name** dialog box, type the name of your new organization. Remember that after you type a name for your Exchange 2000 organization, you cannot change that name.
10. In the **Exchange 2000 Administrator Account** dialog box, type the name of the user or group who is responsible for installing Exchange 2000 Server. This account must be a domain account that includes local administrator privileges on the cluster nodes. The account you specify will also have permission to create all levels of Exchange 2000 administrator accounts with the Exchange delegation wizard. Click **Next**.
11. In the **Completion** dialog box, click **Finish**.

### **Task 2: Run DomainPrep**

You must run DomainPrep in each Windows 2000 domain that you want to install Exchange 2000 Server in. Before you can run the DomainPrep process, replication of the schema updates by the ForestPrep process must finish.

**Note** Running DomainPrep is required only if you are installing Exchange 2000 Server for the first time in your domain. If you have already installed Exchange 2000 Server in your domain, running DomainPrep is not necessary and you may skip to "Step 2: Install Exchange 2000 Server on Each Node" later in this task.

1. Log on to any computer in the domain to which you want to run DomainPrep. The account you use must be a member of the Domain Admins Windows 2000 security group.
2. Click **Start**, and then click **Run**.

3. In the **Open** box, type **CD\_drive\_letter\setup\i386\setup /domainprep**, where *CD\_drive\_letter* is the letter of your CD-ROM drive.
4. In the **Welcome** dialog box, click **Next**.
5. In the **End User License Agreement** (EULA) dialog box, read the EULA, and if you agree to the terms, click **I agree**, and then click **Next**.
6. In the **Product Identification** dialog box, type the 25-digit CD key found on the back of the product CD-ROM, and then click **Next**.
7. In the **Component Selection** dialog box, verify that the action next to Microsoft Exchange 2000 is **DomainPrep**, and then click **Next**.
8. In the **Completion** dialog box, click **Finish**.

## **Step 2: Install Exchange 2000 Server on Each Node**

After you extend the schema with ForestPrep and prepare the domain with DomainPrep, you are ready to install Exchange 2000 on the first cluster node.

### **Task 1: Ensure Cluster Service is Running on Each Node**

Cluster Service must be installed and running in a cluster node to successfully install the cluster-aware version of Exchange 2000.

1. Log on to any node in your Exchange 2000 cluster.
2. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**.
3. In **Cluster Administrator**, in the console tree, select the cluster name under the root container.
4. In the details pane, ensure the **State** of your cluster nodes are all **Up**.

### **Task 2: Install Microsoft Distributed Transaction Coordinator (MSDTC)**

Before installing Exchange 2000, you must first install MSDTC on every node of the cluster.

1. Log on to the node of the cluster to which you want to install Exchange 2000 using a domain account.
2. Click **Start**, and then click **Run**.
3. Type **cmd**, and then click **OK**.
4. At the **Command Prompt**, type **Comclust.exe**, and then click **Enter**.
5. Repeat steps 1-4 of this procedure on all other nodes of the cluster.
6. To verify the installation, in **Cluster Administrator** verify that the MSDTC resource appears in the group **Cluster Group** and is online.

**Note** Comclust.exe may attempt to install an additional component if load balancing is enabled. This component is not needed for the Exchange installation.

### **Task 3: Run Exchange 2000 Setup**

In this task, you are installing the cluster-aware version of Exchange 2000 to each node. Before installing Exchange 2000 on a node, it is recommended that you move all cluster resources owned by the node to another node.

**Note** Install Exchange 2000 completely on one node before you install it on another node.

1. Log on to the node of the cluster to which you want to install Exchange 2000 using the account name you specified in the **Exchange 2000 Administrative Account** dialog box in Step 10 of "Task 1: Run ForestPrep" earlier in this document.
2. Click **Start**, and then click **Run**.
3. In the **Open** box, type **CD\_drive\_letter\setup\i386\setup** where *CD\_drive\_letter* is the letter of your CD-ROM drive.
4. In the **Welcome** dialog box, click **Next**.
5. In the **End User License Agreement** (EULA) dialog box, read the EULA, and if you agree to the terms, click **I agree**, and then click **Next**.
6. In the **Product Identification** dialog box, type the 25-digit CD key found on the back of the product CD-ROM, and then click **Next**.
7. In the **Component Selection** dialog box, verify that the action next to Microsoft Exchange 2000 is **Typical**.
8. To change the installation location of the Exchange 2000 program files, click **Change Folder**. For information about available drives and their corresponding available disk space, click **Disk Information**. By default, the Exchange 2000 program files are installed on the Windows 2000 boot drive. For example, if your Windows 2000 boot files are on drive C, Exchange is installed to C:\Program Files\Exchsrvr. Click **Next**.
9. In the **Licensing Agreement** dialog box, read the per-seat licensing agreement. To continue, click **I agree**, and then click **Next**.
10. In the **Component Summary** dialog box, verify your installation selections, and then click **Next**.
11. When you install Exchange 2000 in a cluster node, Setup prompts you with a dialog box (see Figure 6) that indicates that the cluster-aware version of Exchange 2000 will be installed. In this dialog box, click **OK** to proceed with Setup.



**Figure 6** Cluster-aware version dialog box

12. In the dialog box that reminds you to restart your computer after the installation, click **OK**.
13. In the **Completion** dialog box, click **Finish**.

14. When Exchange 2000 Setup finishes on the first node, restart your computer.
15. Install Exchange 2000 Service Pack 3 (SP3) after restarting your computer.
16. Repeat "Task 3: Run Exchange 2000 Setup," on the second node of the cluster.

**Task 4: Grant the Cluster Administrator Service Account Exchange Full Administrator Permissions.**

After you install Exchange, you must grant the Cluster Service account Exchange Full Administrator privileges. This procedure only needs to be performed once.

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In **Exchange System Manager**, in the console tree, right-click your Exchange organization name, and then click **Delegate control**.
3. In the **Exchange Administration Delegation Wizard** dialog box, click **Next**.
4. In the **Users or Groups** page of the wizard, click **Add**.
5. In the **Delegate Control Box** dialog box, click **Browse**.
6. In **Select Users, Computers or Groups**, specify the Cluster Service account object, and then click **OK**.
7. In the **Delegate Control** dialog box, in **Role**, click **Exchange Full Administrator**, and then click **OK**.
8. Click **Next**, and then click **Finish**.

**Step 3: Create the Exchange Virtual Servers**

The final step in configuring Exchange 2000 on a Windows 2000 cluster is to create the Exchange Virtual Servers. The number of Exchange Virtual Servers you need to create depends on whether you are creating an active/passive or active/active Exchange 2000 cluster.

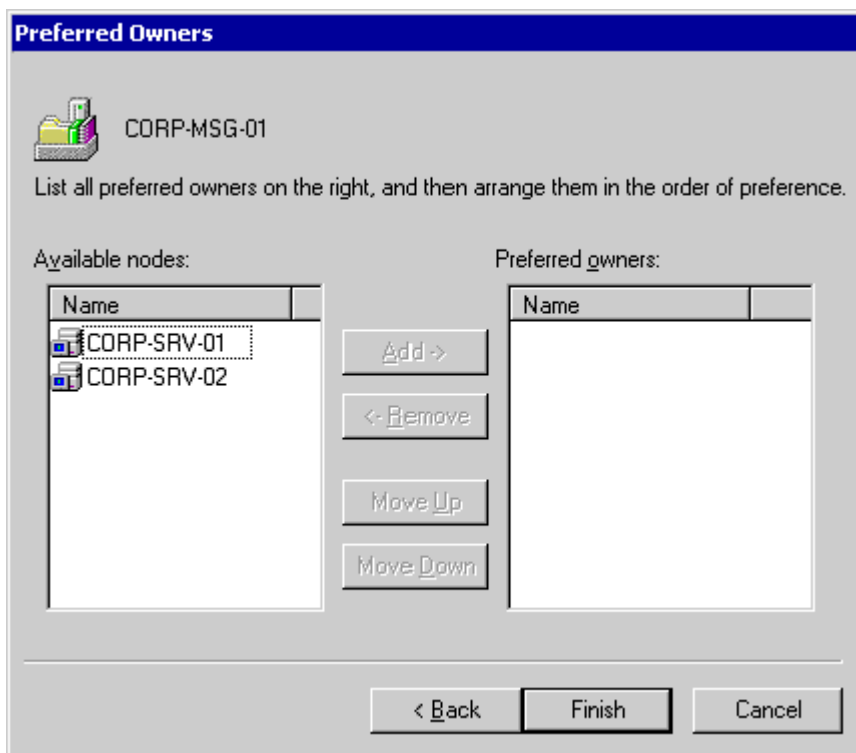
- If you are setting up a two-node active/passive Exchange 2000 cluster, you will be setting up one Exchange Virtual Server. You only need to perform this step once.
- If you are setting up a two-node active/active Exchange 2000 Server cluster, you will be setting up two or more Exchange Virtual Servers. You must repeat this step for each Exchange Virtual Server.

Each Exchange Virtual Server consists of a static IP address, a unique Network Name, a shared physical disk, and an Exchange System Attendant resource.

**Task 1: Create an Exchange Virtual Server**

1. On the first node of the cluster, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Cluster Administrator**. If prompted to specify a cluster, type the cluster name, or browse and select the cluster in which you want to create an Exchange Virtual Server.
2. Right-click the **Groups** container, click **New**, and then click **Group**.
3. The New Group Wizard starts. In the **Name** box, type a name for this Exchange Virtual Server, and then click **Next**.

4. In the **Preferred Owners** dialog box (see Figure 7), verify that there is either one or no cluster nodes listed in the **Preferred Owners** box, and then click **Finish**. This new Exchange Virtual Server (cluster group) is displayed under **Groups**.

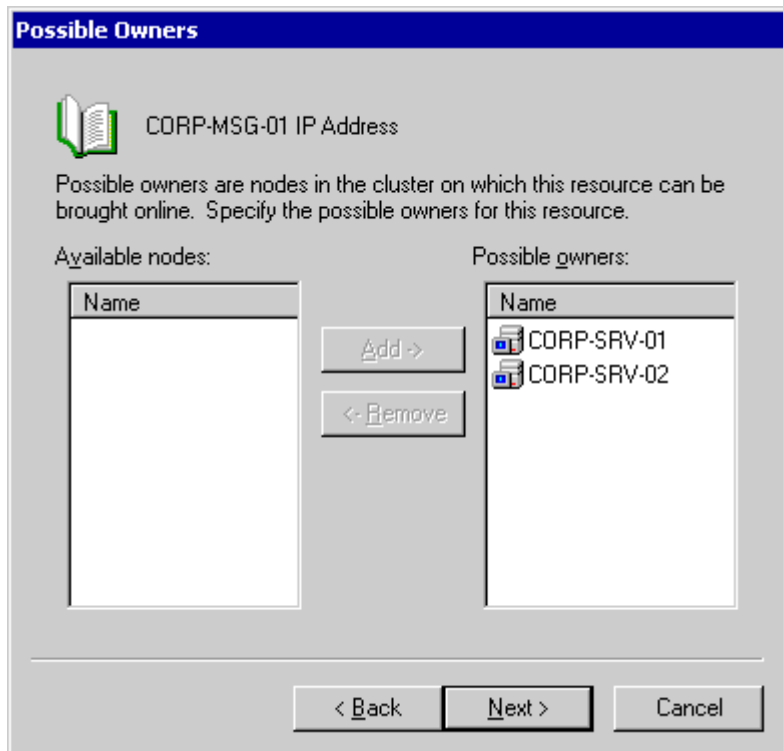


**Figure 7 Preferred Owners** dialog box

**Note** If the **Preferred Owners** dialog box for a group contains both nodes in the cluster, ensure that the order of the list on the other node is opposite. For example, if the preferred owners list on the first node lists CORP-SRV-01 and then CORP-SRV-02, ensure the second node lists CORP-SRV-02 and then CORP-SRV-01.

### **Task 2: Create an IP Address Resource**

1. Right-click the Exchange Virtual Server, point to **New**, and then click **Resource**.
2. The New Resource Wizard starts. In the **New Resource** dialog box, type **EVSName IP Address**, where *EVSName* is the name of your Exchange Virtual Server, in the **Name** box.
3. In the **Resource Type** box, click **IP Address**. Verify that the **Group** box contains the name of your Exchange Virtual Server, and then click **Next**.
4. In the **Possible Owners** dialog box (see Figure 8), verify that all nodes in the cluster that have Exchange installed on them appear in the **Possible owners** box, and then click **Next**.



**Figure 8** Possible Owners dialog box

5. In the **Dependencies** dialog box, verify that no resources appear in the **Resource dependencies** box, and then click **Next**.
6. In the **TCP/IP Address Parameters** dialog box, in **Address**, type the static IP address of the Exchange Virtual Server.

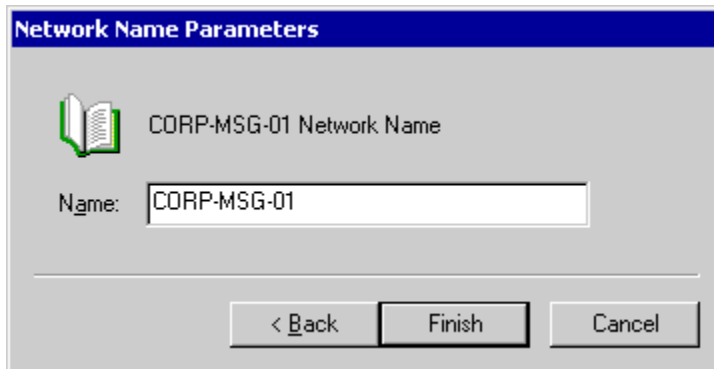
**Note** It is highly recommended that the Exchange cluster have its own dedicated static IP address, separate from all other resources (including the quorum) that are defined in Cluster Administrator.

7. In **Subnet mask**, verify that the subnet mask for the Exchange Virtual Server is correct.
8. In **Network**, verify that the **<Public Network Connection Name>** is selected, and then click **Finish**.

### **Task 3: Create a Network Name Resource**

1. Right-click the Exchange Virtual Server, point to **New**, and then click **Resource**.
2. New Resource Wizard starts. In the **New Resource** dialog box, type **EVSName Network Name**, where *EVSName* is the name of your Exchange Virtual Server, in the **Name** box.
3. In the **Resource Type** box, click **Network Name**, and then click **Next**.
4. In the **Possible Owners** dialog box (see Figure 8), verify that both nodes appear in the **Possible owners** box, and then click **Next**.
5. In the **Dependencies** dialog box, click the **IP Address** resource for this Exchange Virtual Server in the **Available resources** box, and then click **Add**. Click **Next**.

6. In the **Network Name Parameters** dialog box (see Figure 9), type a name for the Exchange Virtual Server in the **Name** box. This name is the network name that identifies the Exchange Virtual Server on your network.



**Figure 9: Network Name Parameters** dialog box

7. Click **Finish**.

#### ***Task 4: Add a Disk Resource to the Exchange Virtual Server***

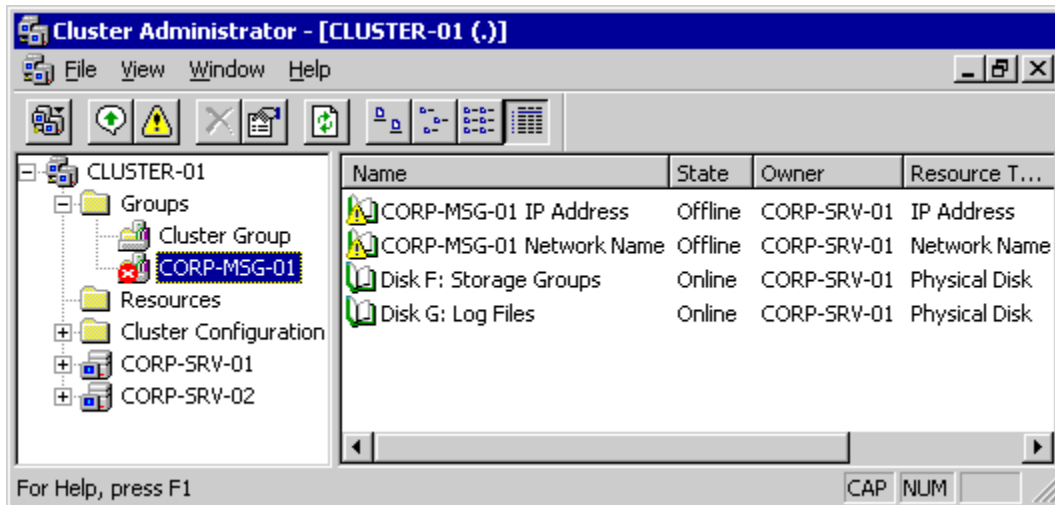
Use this procedure to add a disk resource. You must repeat this task for each disk that you want to associate with the Exchange Virtual Server.

**Note** If the resource you want to add already exists, move the disk resource. If the disk resource you want to add does not yet exist, you must first create a new disk resource.

#### **To move an existing disk resource to an Exchange Virtual Server**

To move an existing disk resource to Exchange Virtual Server, move the disk to the cluster group by doing the following.

1. In **Cluster Administrator**, click the group that contains the physical disk resource you want to move to the Exchange Virtual Server. The node on which you create the Exchange Virtual Server must own this group. If this is not the case, first move the group to this node. You may move the group back to the original node after the move.
2. Drag the physical disk resource to the Exchange Virtual Server. After moving the disk resource, it appears as a resource of the Exchange Virtual Server (see Figure 10).



**Figure 10** Exchange Virtual Server after moving two physical disk resources to it

### To create a new disk resource

**Note** To prevent corruption of your disk, consult the clustering documentation in Windows 2000 Help before connecting a disk to a shared bus.

1. Right-click the Exchange Virtual Server, point to **New**, and then click **Resource**.
2. The **New Resource Wizard** starts. In the **New Resource** dialog box, type **Disk <drive letter>**, where *drive letter* is the logical drive on this disk. You should use a descriptive name, for example "Disk G: Log Files."
3. In the **Resource Type** box, click **Physical Disk**, and then click **Next**.
4. In the **Possible Owners** dialog box (see Figure 8), verify that both nodes appear in the **Possible owners** box, and then click **Next**.
5. In the **Dependencies** dialog box, verify that no resources appear in the **Resource dependencies** box, and then click **Next**.
6. In the **Disk Parameters** dialog box, in **Disk**, select the disk you want. If the disk does not appear here, it means that either another group already has a resource for it, or it was not successfully installed.
7. Click **Finish**. The disk resource appears as a resource of the Exchange Virtual Server (see Figure 10).

### Task 5: Create an Exchange 2000 System Attendant Resource

1. In **Cluster Administrator**, in the console tree, right-click the Exchange Virtual Server, and click **Bring Online** (see Figure 10).
2. Right-click the Exchange Virtual Server, point to **New**, and then click **Resource**.
3. The **New Resource Wizard** starts. In the **New Resource** dialog box, in **Name**, type **Exchange System Attendant - (<EVSName>)**, where *EVSName* is the name of your Exchange Virtual Server, in the **Name** box.
4. In the **Resource Type** box, click **Microsoft Exchange System Attendant**, and then click **Next**.

5. In the **Possible Owners** dialog box (see Figure 8), verify that both nodes appear in the **Possible owners** box, and then click **Next**.
6. In the **Dependencies** dialog box, select both the **Network Name** and **Physical Disk** resources for this Exchange Virtual Server in the **Available resources** box, and then click **Add**. Click **Next**.
7. In the **Data Directory** dialog box, in **Enter path to the data directory**, verify the data directory location. You must verify that this location points to the shared physical disk resource assigned to this Exchange Virtual Server. Exchange will use the drive you select in this step for storing the transaction log files, the default public store files, and the mailbox store files (pub1.edb, pub1.stm, priv1.edb, and priv1.stm). Click **Finish**.
8. Right-click the Exchange Virtual Server, and then click **Bring Online**.

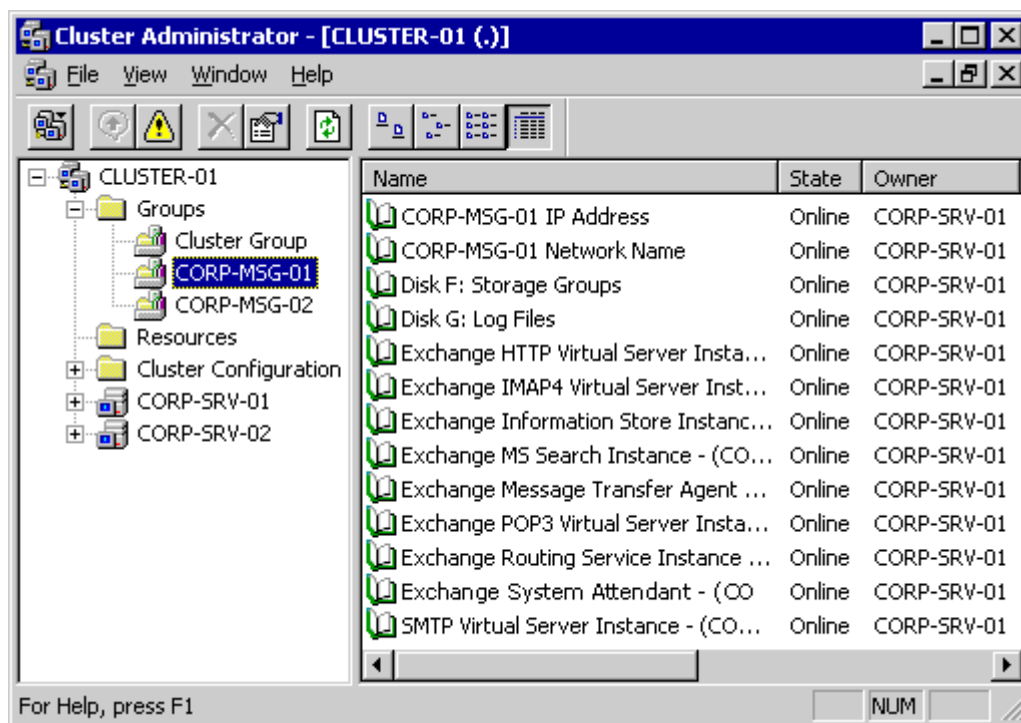
**Note** Due to directory replication latency, all resources may not come online in your first attempt. In this case, wait for replication to occur and bring the resources online again. To add resources to the dependencies list when creating the Exchange System Attendant resource, ensure that the resources you want to add are online before you add them.

After you successfully create the Exchange System Attendant resource, Exchange System Attendant automatically creates all the other resources for the Exchange Virtual Server (see Figure 11).

The Exchange System Attendant resource creates the following resources:

- Exchange Information Store Instance
- Exchange Message Transfer Agent Instance
- Exchange Routing Service Instance
- SMTP Virtual Server Instance
- Exchange HTTP Virtual Service Instance
- Exchange IMAP4 Virtual Server Instance
- Exchange POP3 Virtual Server Instance
- Exchange MS Search Instance

**Note** The Message Transfer Agent Instance resource is only created in the first Exchange Virtual Server added to a cluster. All Exchange Virtual Servers in the cluster share the single Message Transfer Agent Instance resource.



**Figure 11** An active/active Exchange 2000 Server cluster

**Task 6: If You are Creating an Active/Active Cluster, repeat Step 3 for the Second Exchange Virtual Server in the Cluster**

Repeat "Step 3: Create the Exchange Virtual Servers" for each Exchange Virtual Server you want to create. For example, if you are creating a two-node active/active cluster, you must repeat this step. If you are creating a two-node active/passive Exchange 2000 cluster, do not repeat this step.

## Set Up a Four-Node Exchange 2000 Server Cluster

The procedures for setting up a four-node Exchange 2000 Server cluster are very similar to the procedures in the section "Set Up a Two-Node Exchange 2000 Server Cluster" earlier in this document. Use the logic presented in the "Set Up a Two-Node Exchange 2000 Server Cluster" section to help you set up a four-node cluster.

**Note** Your servers must be running Windows 2000 Datacenter Server for four-node clustering.

Before deploying a four-node Exchange 2000 Server cluster, ensure that your four-node Windows 2000 cluster is properly created. For step-by-step instructions on how to create a four-node Windows 2000 Cluster, see the technical paper *Step-by-Step Guide to Installing Cluster Service* at <http://go.microsoft.com/fwlink/?LinkId=266>.

## Cluster Settings

Various cluster-specific settings affect the functionality of Exchange 2000 Server. After you install Exchange 2000 in your cluster nodes and create your Exchange Virtual Servers, you can customize your cluster configuration settings.

## Understanding Exchange Virtual Server and Exchange Resource Settings

All of these settings are either Exchange Virtual Server or Exchange resource specific. Use the Cluster Administrator to change these settings.

### Exchange Virtual Server Configuration Settings

To access the Exchange Virtual Server properties dialog box of an Exchange Virtual Server in a cluster group, in **Cluster Administrator**, in the console tree, right-click the Exchange Virtual Server you want to configure, and then click **Properties**.

On the **General** tab of the Exchange Virtual Server, you can configure the following setting:

- **Preferred Owners.** Use this setting to instruct Cluster Service to assign Exchange Virtual Servers to specific cluster nodes. When a node comes online after being offline, the Exchange Virtual Server attempts to fail back to the preferred node. For example, in a four-node, 3 active/1 passive configuration with three Exchange Virtual Servers, it is recommended that you set the preferred owner setting of each Exchange Virtual Server to a different cluster node giving each node a preferred Exchange Virtual Server.

On the **Failback** tab of the Exchange Virtual Server, you can configure the following setting:

- **Failback** The failback setting specifies whether the group will fail back when a failed node returns to operation. In the previous 3 active/1 passive example with three Exchange Virtual Servers, you have the following two choices in regard to failback:
  - You can choose whether an Exchange Virtual Server should fail back to the preferred node after a failure (when a group is moved from the preferred active node to the passive node) and how long it should wait until it fails back.
  - You can choose not to have the Exchange Virtual Server fail back, which would require administrator intervention to move the Exchange Virtual Server back to the original, preferred node.

The second method is the preferred method because the administrator has control over the time the failback occurs and can coordinate this to give users the least amount of downtime.

### Exchange Resource Configuration Settings

To access the properties of an Exchange cluster resource properties dialog box, in **Cluster Administrator**, in the console tree, click the Exchange Virtual Server that contains the resource you want to configure. In the details pane, right-click the resource you want to configure, and then click **Properties**.

On the **General** tab of a cluster resource's properties, you can configure the following setting:

- **Possible Owners** This cluster resource setting allows the administrator to limit which nodes a resource can run on. Setting this option to allow multiple nodes allows failover to occur. Limiting this option to allow only one possible owner disables

failover. This setting is an effective way to control Exchange Virtual Server failover scenarios.

**Note** It is necessary to set all resources in the Exchange Virtual Server with the same possible owners setting.

On the **Advanced** tab of a cluster resource's properties, you can configure the following settings:

- **Do not restart/Restart** This setting configures the resource to restart or not restart after a resource failure. Select the **Affect the group** check box to configure the resource to cause the Exchange Virtual Server to fail over to a different node. Use the **Threshold** and **Period** settings to determine when this move group should occur. The number of resource failures (threshold) that occur in a certain amount of time (period) determine when the resource will cause the entire group to fail over. It is strongly recommended that you not set the **Do Not Restart** option.
- **Looks Alive Poll Interval** This feature is not implemented for Exchange 2000 Server cluster resources.
- **Is Alive Poll Interval** Cluster Service performs a comprehensive check to see if the cluster resource is active or not. All Exchange 2000 Server cluster resources use an Is Alive poll of 10 seconds. This value cannot be changed.
- **Pending timeout** This setting provides a space for you to type the amount of time, in seconds, that a resource in a pending state (online pending or offline pending) is given to resolve its status before Cluster Service puts the resource in the failed state.

By default, this timeout is 180 seconds or 3 minutes. Depending on your **Do not restart/Restart** settings, the resource will stay failed or will initiate a restart. With the exception of the Exchange store instance, all Exchange 2000 and Windows 2000 cluster resources can go offline and online during this period. The Exchange Information Store Instance resource may need more time to go offline or online. The amount of time is based on the state of the Exchange database and logs at the time of startup or shutdown. The load on the server as well as the number of storage groups and databases are also factors. Thus, it is often necessary to increase the pending timeout value for the Exchange Information Store Instance resource.

Depending on the load and configuration of the servers, a timeout set between 5 and 10 minutes is usually adequate. It is important to note that, if the Exchange Information Store resource were to stop responding during an offline or online pending state, because of the increased pending timeout setting, it may take Cluster Service longer to notice that the resource is in a bad state and fail it over to another group.

It is also important to note that, if Cluster Service fails over an Exchange Information Store resource because of a pending timeout failure, the Extensible Storage Engine (ESE) needs to run through recovery on the databases and may take much longer to restart the resource.

## Exchange Logging

After you install Exchange 2000 on your cluster nodes, and create your Exchange Virtual Server, you may want to configure Exchange logging. Although it is helpful to enable

Exchange logging when troubleshooting message flow issues, it is generally not recommended that you enable logging at all times. This is because logging reduces e-mail performance. For more information about monitoring performance of your Exchange cluster, see "Monitoring Performance" later in this document.

If you enable logging when using Exchange clustering, you should perform the following tasks in this section.

**Important** It is only necessary to perform these procedures if you need to enable Exchange logging.

### **Disabling Microsoft Exchange MTA Stack Service Monitoring**

The default monitoring configuration on an Exchange 2000 server includes the Microsoft Exchange MTA Stacks service. Because the MTA in a cluster environment only runs on one of the physical computers, monitoring reports that the computer, without the MTA running on it, is in an error state. This, in turn, can cause problems if Exchange 2000 is installed on a cluster computer with two or more Exchange Virtual Servers. You should perform this task on the second Exchange Virtual Server (and if applicable, any other additional Exchange Virtual Servers) of a cluster. You do not need to perform this task on the first Exchange Virtual Server of a cluster.

#### **To disable monitoring of the Microsoft Exchange MTA**

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In **Exchange System Manager**, in the console tree, expand **Servers**, right-click the server on which you just installed Exchange, and then click **Properties**.
3. In the **<Server Name> Properties** dialog box, click the **Monitoring** tab.
4. On the **Monitoring** tab, select **Default Microsoft Exchange Services** from the list of services, and then click **Details**.
5. In the **Default Microsoft Exchange Services** dialog box, select **Microsoft Exchange MTA Stacks**, and then click **Remove**.
6. Click **OK** twice.

### **Specifying Exchange to Log SMTP to a Shared Disk in Your Cluster**

The log files created by SMTP logging can be used to gather statistical data on server usage, and can be a useful tool for administrators. SMTP logging is not enabled by default because it reduces Exchange performance. When enabled, Exchange creates SMTP log files on the system drive of the local computer (for example, C:\Winnt\System32\Logfiles, where C is the location of your Windows 2000 installation). To reliably configure SMTP logging, you must specify a folder on a shared disk.

#### **To enable SMTP logging and to specify that those log files are logged to a shared disk**

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In **Exchange System Manager**, in the console tree, expand **Servers**, and then expand the server to which you want to enable SMTP logging.

3. In the console tree, expand **Protocols**, and then expand **SMTP**.
4. In the console tree, right-click **Default SMTP Virtual Server**, and then click **Properties**.
5. In the **Default SMTP Virtual Server Properties** dialog box, on the **General** tab, click **Enable logging**, and then click **Properties**.
6. In the **Extended Logging Properties** dialog box, on the **General Properties** tab, in **Log file directory**, change the SMTP log file location to a folder on a shared disk.
7. Click **OK** twice.

## Configure a Clustered Back-End Server

If you have any Exchange 2000 servers in your organization configured as front-end servers, it is necessary to modify settings using both Cluster Administrator and Exchange System Manager.

To configure a clustered back-end server, you must map each front-end server to both nodes of your cluster so that either node can accept proxy requests from any front-end server in your organization. Proxy requests are requests for messaging services from client computers running Outlook Web Access, POP3, or IMAP that are sent to the cluster through the front-end servers. All communication between front-end and back-end server goes through TCP port 80, regardless of the port used for communication between client and front-end.

For more information about configuring front-end and back-end servers running Exchange 2000, see the technical article *Exchange 2000 Front-End Back-End Topology* at <http://go.microsoft.com/fwlink/?LinkId=4721>.

The following steps describe how to configure the HTTP virtual server resources in an Exchange Virtual Server to function as a back-end server.

1. Create the HTTP virtual servers in Exchange System Manager.
2. Create virtual directories to match the directories configured on the front-end server.
3. Add new HTTP virtual server resources to the Exchange Virtual Server.

**Note** In this section and throughout this document, Exchange Virtual Server refers to the Exchange Virtual Server in the cluster.

### Step 1: Create the HTTP Virtual Servers in Exchange System Manager


When you create an Exchange Virtual Server, during the installation of the System Attendant resource, Exchange installs the HTTP virtual server resources as well as other Exchange resources. These HTTP virtual server resources are configured with the host header set to the network name of the Exchange Virtual Server. To configure a front-end server to use a clustered back-end server, additional HTTP virtual servers must be created on each Exchange Virtual Server that is a part of the clustered back-end servers.

You must create one Exchange HTTP virtual server for each front-end namespace. For example, if contoso.com hosts Exchange 2000 Server for tailspintoys.com and wingtiptoys.com, you will have three virtual servers—the default virtual server, a virtual server for tailspintoys.com, and a virtual server for wingtiptoys.com. This configuration

allows for maximum flexibility in determining which resources are available to each hosted company.

You must perform the following steps to configure the Exchange Virtual Server to function as a back-end server. These steps must be repeated for each Exchange Virtual Server in the cluster:

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In **Exchange System Manager**, in the console tree, expand **Servers**, and then expand the server that you want to configure as a back-end server.
3. In the console tree, expand **Protocols**, and then expand **HTTP**.
4. In the console tree, right-click **HTTP**, point to **New**, and then click **HTTP** virtual Server.
5. In **Name** field, type the name of your front-end server.
6. Next to the **IP Address** list box, click **Advanced**.
7. In the **Advanced** dialog box, click **Add**.
8. In the **Identification** dialog box (see figure 12), in **IP address**, select the IP address of this Exchange Virtual Server (the back-end server). This IP address must match the IP address resource value you previously configured for this resource.



**Figure 12** Back-end server **Identification** dialog box

9. In the **Host Name dialog** box, type the host header of the front-end server. This is the name with which the front-end server is accessed by clients. The host header for the Exchange Virtual Server must map to the host header on the front-end server.

**Note** Client requests to the front-end server use a specific host, such as `http://mail.contoso.com`. A virtual server on the front-end must have the "mail.contoso.com" host header configured. The front-end server then proxies the request to the back-end server, which must also have the host header configured on a virtual server.

10. Verify that **TCP port** is set to 80, and then click **OK**.
11. In the **Advanced** dialog box, if you want to add an additional identity, click **Add**, and follow Steps 8 through 10 again.

**Note** Consider adding several identities to the virtual server that list all the ways that a user might access the front-end server. For example, if a front-end server is used both internally and externally, consider listing both a host name and a fully qualified domain name, such as "mail" for internal access and "mail.contoso.com" for external access.

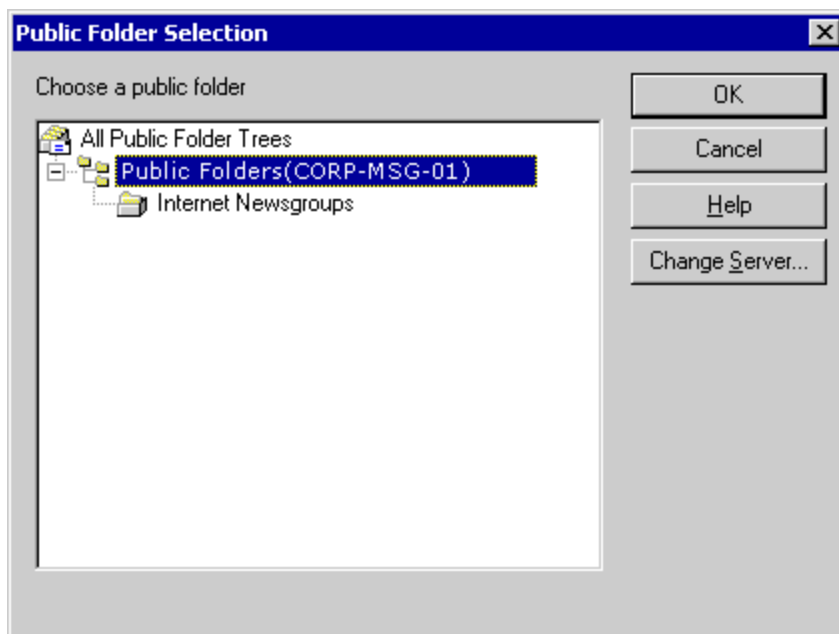
12. In the **Advanced** dialog box, click **OK** twice to create the new HTTP virtual server.

### **Step 2: Create Virtual Directories to Match the Directories Configured on the Front-End Server**

After you create the HTTP Virtual Server in Exchange System Manager, you must add virtual directories to match those configured on the front-end server in Exchange System Manager. In a typical Exchange installation, you will have the virtual directories **Exchange** and **Public**. In System Manager, virtual directories of HTTP Virtual Servers display as child objects of the HTTP Virtual Server.

Do the following to create virtual directories **Exchange** and **Public** on the back-end server:

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In **Exchange System Manager**, in the console tree, expand **Servers**, and then expand the server that you want to configure as a back-end server.
3. In the console tree, expand **Protocols**, and then expand **HTTP**.
4. In the console tree, under **HTTP**, right-click **<HTTP Virtual Server Name>** (where *HTTP Virtual Server Name* is the name of the HTTP virtual server you created in "Step 1: Create the HTTP Virtual Servers in System Manager" earlier in this document), point to **New**, and then click **Virtual Directory**.
5. In the **Properties** dialog box, in **Name**, type **Exchange**.
6. Under **Exchange Path**, the **Mailboxes for** option is selected by default. Keep this setting because the virtual directory Exchange users will use it to access their Exchange mailboxes. Click **OK** to create the first virtual directory.
7. In the console tree, under **HTTP**, right-click **<HTTP Virtual Server Name>** (where *HTTP Virtual Server Name* is the name of the HTTP virtual server you created in "Step 1: Create the HTTP Virtual Servers in Exchange System Manager" earlier in this document), point to **New**, and then click **Virtual Directory**,
8. In the **Properties** dialog box, in **Name**, type **Public**.
9. Under **Exchange Path**, click **Public folder**, and then click **Modify**.
10. In the **Public Folder Selection** dialog box, double-click **Public Folders**. After a few seconds, Exchange resolves the public folder's server name and appends it to the name of the Public Folders container (see Figure 13).



**Figure 13** Public Folder Selection dialog box after expanding **Public Folders**

11. To close the **Public Folder Selection** dialog box, click **OK**.
12. To create the second virtual directory, click **OK**.
13. If you have any additional virtual directories configured on your front-end server, you must also create those virtual directories.

For any virtual directories that point to mailboxes, ensure the SMTP domain selected on the properties of the virtual directory matches the SMTP domain of users who will be using that front-end server. If the proper domain is not selected, users of that domain will not be able to log on using that virtual server. The list of domains is pulled from the domains of the SMTP addresses in the Exchange organization's recipient policies. If you have more than one recipient policy for the same domain, you will see duplicates. In this case, it does not matter which one you choose.

For additional information about performing these steps, see Exchange 2000 Help.

### **Step 3: Add New HTTP Virtual Server Resources to the Exchange Virtual Server**

After each HTTP virtual server is created in Exchange System Manager, it must be added to the Exchange Virtual Server in the Cluster Administrator so that Cluster Service can manage it. Do the following to add a new HTTP virtual server resource to the Exchange Virtual Server.

**Note** You must perform this step for each Exchange Virtual Server to which you have added a new HTTP virtual server.

1. From Cluster Administrator, right-click the Exchange Virtual Server, point to **New**, and then click **Resource**.
2. In the **New Resource** dialog box, in **Name**, type **Exchange HTTP Virtual Server - (<EVSNamE>)**, where *EVSNamE* is the name of the front-end server.

3. In the **Resource Type** drop-down list, click **Microsoft Exchange HTTP Server Instance**. Verify that the **Group** drop-down list contains the name of your Exchange Virtual Server, and then click **Next**.
4. In the **Possible Owners** dialog box (see Figure 8), verify that all nodes appear in the **Possible owners** list box, and then click **Next**.
5. In the **Dependencies** dialog box, add the **Exchange Information Store Instance** to the **Resource dependencies** list box, and then click **Next**.
6. In the **Virtual Server Instance** dialog box, in the **Server Instance** drop-down list, select the newly created HTTP virtual server for the resource, and then click **Finish**.
7. To bring the newly created resource online, you must bring the Exchange Virtual Server offline, restart IIS, and then bring your Exchange Virtual Server back online. Do the following to complete this procedure:
  1. From **Cluster Administrator**, right-click the Exchange Virtual Server, and then click **Take Offline**.
  2. To restart IIS, click **Start**, and then click **Run**.
  3. In **Run**, type **cmd**, and then click **Enter**.
  4. To restart IIS, at the command prompt, type **Iisreset.exe**.
  5. After restarting IIS, in **Cluster Administrator**, in the console tree, right-click the Exchange Virtual Server, and then click **Bring Online**.
8. Repeat the three steps in this section for the other Exchange Virtual Servers in the cluster that you want to configure as back-end servers.

## Removing Exchange 2000 from a Cluster

Use the following procedures either to remove a single Exchange Virtual Server from a cluster (for example, changing from an active/active cluster to an active/passive cluster) or to remove an Exchange 2000 Server cluster completely.

In each case, ensure you back up critical data and secure resources hosted by the computer running Exchange 2000.

**Important** Deleting components of an Exchange Virtual Server without removing the entire Exchange Virtual Server can cause interruptions in mail flow. As a result, following this procedure to remove an Exchange Virtual Server from a cluster is recommended.

## Removing an Exchange Virtual Server from a Cluster

This procedure removes one Exchange Virtual Server from a cluster. Use this procedure either as the first step in removing the Exchange 2000 Server cluster, or for changing your Exchange 2000 Server cluster from a two-node active/active to a two-node active/passive cluster.

**Important** If your cluster includes more than one Exchange Virtual Server, do not remove the Exchange Virtual Server that owns the Message Transfer Agent Instance resource until you remove all other Exchange Virtual Servers in the cluster. The first Exchange Virtual Server created in a cluster owns the Message Transfer Agent

Instance resource. All other Exchange Virtual Servers are dependent on this resource.

If you need to remove the first Exchange Virtual Server in a cluster, you must first remove all other Exchange Virtual Servers in that cluster and then finally remove the first Exchange Virtual Server.

To remove an Exchange Virtual Server, you must perform the following five tasks:

1. Move all mailboxes and public folder content to another Exchange Virtual Server.
2. Take the System Attendant resource offline.
3. Delete the System Attendant resource. (This step removes all the resources that comprise an Exchange Virtual Server from the cluster resource group.)
4. Ensure that the Exchange Virtual Server server object is deleted from Active Directory.
5. Delete remaining cluster resources.

**Important** Before deleting an Exchange Virtual Server, you should determine if that server is used as a bridgehead server by a connector. If it is, you should designate another server as the bridgehead server.

### **Task 1: Move All Mailboxes and Public Folder Content to Another Exchange Virtual Server**

The first task you must complete before deleting an Exchange Virtual Server is to move any mailboxes that reside on the Exchange Virtual Server you want to remove. You can move the mailboxes to other servers in your Exchange organization. If you do not move the mailboxes, they will be deleted when you remove the Exchange Virtual Server.

To move mailboxes from one server (source) to the final destination server (target), use Active Directory Users and Computers. Right-click the users, click **Exchange Tasks**, and then click **Move Mailbox**. For more information about moving mailboxes, see Microsoft Knowledge Base article [Q224975](#) "XADM: Differences Between Move Mailbox Methods." For information about moving a large number of mailboxes, see Microsoft Knowledge Base article [Q297393](#) "HOW TO: Programmatically Move an Exchange 2000 Mailbox Using CDOEXM in Visual C++."

If this Exchange Virtual Server stores public folder content, you must also move the public folder content to another server. For information about moving public folders, see Microsoft Knowledge Base article [Q288150](#) "XADM: How to Rehome Public Folders in Exchange 2000."

### **Task 2: Take System Attendant Resource Offline**

1. From **Cluster Administrator**, select the Exchange Virtual Server you want to remove.
2. In the details pane, right-click **System Attendant** resource, and then click **Take Offline**. Exchange System Attendant for this Exchange Virtual Server will be taken offline, as well as its dependant resources.

### Task 3: Delete Exchange System Attendant Resource

1. From **Cluster Administrator**, select the Exchange Virtual Server you want to remove.
2. In the details pane, right-click **System Attendant** resource, and then click **Delete**.
3. Cluster Administrator also deletes the dependant resources for the Exchange System Attendant; in the **Delete Resources** dialog box, click **Yes** to delete the dependant resources. The Exchange System Attendant resource for this Exchange Virtual Server and its dependant resources are deleted. It also removes the Exchange Virtual Server information from Active Directory, leaving the physical disk, the IP address, and Network Name resources.

### Task 4: Ensure the Exchange Virtual Server Server Object is Deleted from Active Directory

1. On the **Start** menu, point to **Programs**, point to **Microsoft Exchange**, and then click **System Manager**.
2. In Exchange System Manager, in the console tree, expand **Servers**.
3. In the console tree, select the **Servers** node and verify that the server object for the Exchange Virtual Server does not appear in the details pane.
4. If the deleted Exchange Virtual Server server object remains, wait for Active Directory replication to occur. Alternatively, you can force replication using the **Replicate Now** option of the Active Directory Sites and Services snap-in.

**Note** If you cannot remove the Exchange Virtual Server object, search the Microsoft Knowledge Base for a resolution or contact Microsoft Product Support Services at <http://go.microsoft.com/fwlink/?LinkId=5967>.

### Task 5: Delete Remaining Cluster Resources

1. From **Cluster Administrator**, select the cluster group that contained the Exchange Virtual Server you just deleted.
2. In the details pane, right-click the **IP Address resource**, and then click **Take Offline**.
3. Right-click the **IP Address resource** again, and then click **Delete**. In the **Delete Resources** dialog box, click **Yes**. Both the **IP Address** resource and the **Network Name** resource will be deleted.
4. Move the **Physical Disk** resource to another group. You can move it by dragging and dropping it into the desired destination group. The destination group must be owned by this node.
5. Delete the cluster group by right-clicking the group in the console tree and selecting delete.

**Important** This Exchange Virtual Server has been deleted. If you want to keep other Exchange Virtual Servers, do not remove the Exchange 2000 Server installation from the node. If you want to completely remove the Exchange 2000 installation, proceed to "Remove an Exchange 2000 Installation from a Cluster" in the next section of this document.

## Remove Exchange 2000 Installation from a Cluster


The following procedure completely removes Exchange 2000 from a cluster. This procedure removes Exchange 2000 installation files from the node. If you want Exchange 2000 to be able to use the node, for example as a passive node, do not uninstall Exchange 2000 from the node. Exchange 2000 cluster configurations, in which only some nodes in a cluster function as Exchange 2000 clusters, are not supported. This means that it may be necessary for you to evict the node from the cluster if other nodes in the cluster will host Exchange Virtual Servers.

**Important** Before you begin, perform the steps in "Removing an Exchange Virtual Server from a Cluster" earlier in this document for every Exchange Virtual Server in the cluster. Alternatively, you can simply evict the cluster node that owns the Exchange Virtual Server you want to remove. If you have already performed this procedure, proceed to the following procedure.

**Note** If the node you are removing owns any important cluster resources, you should move them to another node before proceeding. If you do not move them manually, any remaining resources are automatically moved to other nodes in your cluster when you restart your computer at the end of this procedure.

### To remove the Exchange 2000 Server installation from the node

1. Create a temporary folder (not a subfolder of the Exchange folder).
2. Copy the following files from the Exchange \BIN directory (for example, D:\Program Files\Exchsrvr\BIN, if you originally installed Exchange in the D:\Program Files\Exchsrvr directory), to the temporary folder: Dsaccess.dll, Exchmem.dll, Exproxy.dll, Pttrace.dll.
3. From **Control Panel**, double-click **Add/Remove Programs**.
4. In the **Currently Installed Programs List**, select **Microsoft Exchange 2000**.
5. Click **Change/Remove**.
6. In the **Welcome** dialog box, click **Next**.
7. In the **Component Selection** dialog box, ensure that the action next to Microsoft Exchange 2000 is **Remove**, and then click **Next**.
8. In the **Component Summary** dialog box, verify your installation selections, and then click **Next**.
9. In the **Microsoft Exchange 2000 Wizard** dialog box, click **Yes** if you are removing the last node in the cluster, or click **No** if it is not the last node (see Figure 14). If you remove Exchange from the last node in the cluster, Exchange 2000 Setup removes Exchange cluster resource types from the cluster.



**Figure 14** Removing Exchange 2000 from a cluster

10. In the **Completion** dialog box, click **Finish**.
11. Restart your computer. After restarting, you will see a message box with the title **exmgmt.exe - Unable To Locate DLL** and another message informing you that **At least one service or driver failed during system startup**. This is expected.
12. Copy the DLLs that you copied to a temporary folder in Step 2 back to their original location (for example, to D:\Program Files\Exchsrvr\BIN, if you originally installed Exchange to the D:\Program Files\Exchsrvr directory).
13. Click **Start**, and then click **Run**.
14. Click **Browse**, navigate to the D:\Program Files\Exchsrvr\BIN directory, where D:\Program Files\Exchsrvr is the original location of your Exchange installation, and then click **Exmgmt.exe**.
15. In **Run**, add the parameter **\uninstall**, and then click **Enter**.

**Note** No messages or alerts are displayed by Exmgmt.

## Reliability

Creating Exchange 2000 clusters provides high levels of reliability and allows you to keep your mission-critical applications running in the event of a failure. The following sections discuss reliability in Exchange 2000 clusters.

A common IT industry term for maximum reliability is “five nines,” meaning that a server runs 99.999 percent of the time, which translates into just 5 minutes of downtime per year. Most businesses, however, do not need such stringent uptime requirements.

For the majority of usage scenarios, 99.99 percent uptime is adequate, because this percentage equals less than one hour of downtime per year. The Aberdeen Group found that Windows 2000 servers delivers 99.95 percent uptime right out of the box, before the servers are fully optimized for the environment, and before the IT staff is experienced using the new operating system.

Building Exchange 2000 Server clusters harnesses the reliability of Windows 2000 and allows you to build highly available Exchange 2000 clusters. To read The Aberdeen Group’s report and to find out more about Windows 2000 reliability, visit the Windows 2000 Web site at <http://go.microsoft.com/fwlink/?LinkId=9066>.

## Failover

The failover time for Exchange Virtual Servers is important. To maintain high availability, the failover time must be short. There are two scenarios for failover: planned and unplanned.

### In a planned failover:

- The Exchange administrator requests that the Exchange Virtual Server be moved to another node using Cluster Administrator.
- All resources of the Exchange Virtual Server go offline.
- The resources are moved to the node specified by the Exchange administrator.

- All resources of the Exchange Virtual Server go online.

#### **In an unplanned failover:**

- One resource (or several) of the Exchange Virtual Server fails. The failure is discovered with the next IsAlive check.
- All dependent resources are automatically taken offline by Cluster Service.
- If the failed resource is configured to restart (default setting), Cluster Service will attempt to restart the failed resource and all its dependent resources.
- If the resource fails again:
  - Cluster Administrator will try to restart it again,
  - Or-
  - If the resource is configured to affect the group (default) and the resource has failed a certain number of times (default 3) within a configured time period (default 300 seconds), all resources in the Exchange Virtual Server are taken offline by Cluster Administrator.
- All resources are failed over (moved) to another node in the cluster. If specified, this is the next node in the **Preferred Owners** list.
- Cluster Administrator will attempt to bring all resources of the Exchange Virtual Server online on the new node.
- If the same or another resource fails again on the new node, Cluster Administrator will repeat the steps above, and may need to fail over to yet another node (or back to the original node).
- If the Exchange Virtual Server keeps failing over, Cluster Administrator will fail over the Exchange Virtual Server a maximum number of times (default 10) within a specified time period (default 6 hours). After this time, the Exchange Virtual Server will stay in a failed state.
- If failback is configured (default is turned off), Cluster Administrator will either move the Exchange Virtual Server back to original node immediately when the original node becomes available or at a specified time of day if the original node is available again, depending on the group configuration.

## **Storing Exchange Data**

Exchange stores data in three main locations:

- SMTP queue directory
- .edb and .stm files
- Transaction log files

### **SMTP Queue Directory**

The SMTP queue stores SMTP messages until they are written to a database (public or private, depending on the type of message), or sent to another server or connector.

Typically, messages stored in the SMTP queue are there for a short time. Therefore, your storage solution for the SMTP queue should optimize performance before capacity

and reliability. However, in some situations when downstream processes fail, the SMTP queue could be required to store a large amount of data. For that reason, do not assume that a RAID-0 array is the best storage solution for SMTP queues. Generally, RAID-0 is acceptable only if mail loss is acceptable. RAID-1 is a good solution because it gives some measure of reliability while providing adequate throughput.

### **.edb and .stm Files**

An Exchange database consists of a rich-text .edb file and a native multimedia content .stm file.

The .edb file stores the following items:

- All of the MAPI messages
- Tables used by the Store.exe process to locate all messages
- Checksums of both the .edb and .stm files
- Pointers to the data in the .stm file

The .stm file contains messages that are transmitted with their native Internet content. Because access to these files is generally random, they can be placed on the same disk volume.

As you plan your storage solution for these files, implement a solution that ensures reliability; in other words, RAID-0 is not a recommended option. After reliability, your storage solution is based on a choice between optimizing performance (RAID-1) and optimizing capacity (RAID-5). If possible, use RAID-1 (or RAID-0+1) for these files.

You can store public folders on a RAID-5 array because data in public folders is usually written once and read many times. RAID-5 provides improved read performance.

### **Transaction Log Files**

The transaction log files maintain the state integrity of your .edb and .stm files, which means that the log files represent the data. There is a transaction log file database for each storage group. This file is implemented as a database to increase performance. If a disaster occurs and you have to rebuild your server, use the latest transaction log files to rebuild your databases. If you have the log files and latest backups, you can recover all your data. If you lose your log files, however, the data is lost.

The level of disaster recovery you can achieve relates to the log files (for example, a storage group). For optimal reliability and performance, keep the log files for each storage group on their own drive. As you plan your storage solution for the transaction log files, make integrity and reliability of the utmost importance. Thus, a mirrored RAID-1 (or RAID-0+1) solution is recommended.

## **Performance and Monitoring**

Ensuring that your Exchange 2000 Server clusters perform well involves setup steps and proactive monitoring of your clusters. The following sections provide steps to improve performance, monitor performance, and test the performance of your Exchange 2000 clusters.

## Memory

By default, Exchange 2000 uses all of the physical memory available on your computer; however, you can restrict the amount of memory used by Exchange.

The biggest individual consumer of memory in Exchange 2000 is the Store.exe process. On an active, production Exchange 2000 Server, it is not uncommon to notice that the Store.exe process consumes nearly all of the server memory. Like Exchange Server version 5.5, the Store.exe process uses a unique cache mechanism called Dynamic Buffer Allocation (DBA). This process self-governs how much memory it uses, and balances that with other applications running on the server. If Exchange is the only application running, DBA allocates more memory to itself.

The amount of memory you need in your server depends on the number of databases, size of the databases, and number of transactions. As you create more Exchange databases on the server, your memory requirements increase. Database configuration achieves a certain amount of memory preservation. For example, the first database in a storage group consumes the greatest amount of virtual memory; whereas, if you add a new database to an existing storage group, the memory consumption of this database will be less.

Exchange 2000 can handle up to 20 databases per server (or per cluster node); the total storage space is made up of a maximum of four storage groups and five databases per storage group. Wherever possible, fill out your storage groups to the maximum number of databases before you create a new storage group.

The advantages to filling out your storage group are:

- Reduced memory consumption
- Reduced disk overhead

However, there are a few disadvantages:

- Circular logging can only be controlled at the storage group level and is not recommended.
- Only one backup process can take place in a single storage group at a time. Backing up one database in a storage group will halt the online maintenance of all other databases in the storage group.

## Virtual Memory

Windows 2000 implements a virtual memory system based on a flat (linear) 32-bit address space. Thirty-two bits of address space translates into 4 GB of virtual memory. On most systems, Windows 2000 allocates half this address space (the lower half of the 4-GB virtual address space from x00000000 through x7FFFFFFF) to processes for its unique private storage and the other half (the upper half, addresses x80000000 through xFFFFFFFF) to its own protected operating system memory usage.

It is important to monitor the virtual memory on your Exchange 2000 clusters. For more information about monitoring virtual memory, see the "Monitoring Performance" section later in this document.

For more information about virtual memory, see Windows 2000 Help, *Microsoft Windows 2000 Server Resource Kit*, and *Inside Microsoft Windows 2000*.

## **/3GB Switch**

If your computer running Exchange 2000 has 1 GB of physical memory or more, it is very important to add the **/3GB** switch to the Boot.ini file on the server so that 3 GB of virtual address space is available for user-mode applications. By default, Microsoft Windows 2000 Advanced Server reserves 2 GB of virtual address space for the kernel and allows user mode processes, such as the Exchange 2000 Store.exe process, to use 2 GB of virtual address space.

Virtual address space for a specific process is allocated at startup and increases as more memory is used during run time. It is normal for the actual memory usage of a process to be much less than the address space the process was allocated.

**Important** If your computer running Exchange 2000 has 1 GB of memory or more, it is very important that the Store.exe process does not run out of virtual address space. If it runs out of virtual address space, memory allocation fails, even if there is plenty of physical RAM left, and you must restart the Microsoft Exchange Information Store service.

### Example

A server with 2 GB of physical RAM without the /3GB switch in the Boot.ini file will run out of memory when the Store.exe process's memory usage reaches 2 GB. Windows Task Manager will show that only about 1.5 GB is being used when, in fact, the server is out of memory. The /3GB switch allows processes to allocate up to 3 GB of memory.

## **Exchange 2000 Capacity and Topology Calculator**

Exchange 2000 Capacity and Topology Calculator helps you determine the size of the servers you need for your Exchange 2000 Server clusters. Exchange 2000 Capacity and Topology Calculator can be found on the Microsoft Exchange Web site at <http://go.microsoft.com/fwlink/?LinkId=1716>.

### **Sizing Active/Passive Clusters**

Active/passive clusters are the recommended configuration for Exchange 2000 Server clusters. Windows 2000 Advanced Server supports two-node active/passive clusters, and Windows 2000 Datacenter Server supports two-node, three-node, or four-node active/passive clusters.

Sizing for active/passive clusters can be done using Exchange 2000 Capacity and Topology Calculator, just as you would for stand-alone server deployments.

**Note** It is highly recommended that you test the sizing metrics you determine in a lab using Load Simulator (LoadSim) before deployment. For more information about LoadSim, see the "Microsoft Exchange Load Simulator" section later in this document.

### **Sizing Active/Active Clusters**

Active/active clusters are not a recommended configuration for Exchange 2000 clusters. Exchange 2000 supports two-node active/active clusters only.

When using Exchange 2000 Capacity and Topology Calculator to plan your active/active clusters, there are two very important constraints:

- Ensure that the number of concurrent user connections per node does not exceed 1,900. If you have more than one Exchange Virtual Server per node, ensure that the sum of all concurrent user connections is less than 1900.
- Ensure the average CPU load per server does not exceed 40 percent.

**Note** Before you deploy, you must test the sizing metrics you determine in a lab using LoadSim. For more information about LoadSim, see the “Microsoft Exchange Load Simulator” section later in this document.

After you deploy your cluster, you must do the following:

- Monitor the number of concurrent connections (users) per node. If the number of concurrent users per node exceeds 1,900 for more than 10 minutes, move users off the node.
- Monitor the CPU load for each server in the cluster. If the CPU load exceeds 40 percent (load generated from users) for more than 10 minutes, move users off the server. This load does not include administrative increases in loading, for example, moving users.

## Testing Server Capacity

It is extremely important to test the capacity of your Exchange 2000 Server clusters before you make them available in your organization.

The following list identifies some of the hardware components you need to test:

- Individual computer components such as hard disks and controllers, processors, and RAM
- External components such as routers, bridges, switches, cabling, and connectors

The following are some of the stress tests you need to set up:

- Test cluster performance under heavy network loads
- Test cluster performance under heavy disk input/output (I/O) to the same disk
- Test cluster performance under heavy Exchange services load
- Test cluster performance under large number of simultaneous logon attempts
- Fail over each Exchange Virtual Server at least once to each of the nodes

## Microsoft Exchange Load Simulator

Microsoft Exchange LoadSim allows you to simulate the load of MAPI clients against Exchange 2000 Server. You simulate the load by running LoadSim “tests” on client computers. These tests send messaging requests to the Exchange server causing a load.

Use the output from these tests in the following ways:

- To calculate the client response time for the server configuration under client load
- To calculate the realistic number of users per server
- To identify bottlenecks on the server

For more information about the LoadSim tool or to download LoadSim, see the "Load Simulator" section of *Microsoft Exchange 2000 Server Resource Kit* at <http://go.microsoft.com/fwlink/?LinkId=1710>.

### Exchange Stress and Performance (ESP) Tool

ESP is a highly scalable stress and performance tool for Exchange 2000 Server. It simulates large numbers of client sessions by concurrently accessing one or more protocol services. Scripts control the actions each simulated client takes. The scripts contain the logic for communicating with the server. Test modules (DLLs) then run these scripts. Test modules connect to a server through Internet protocols, calls to application programming interfaces (APIs), or through interfaces like OLE DB.

ESP is modular and extensible and currently provides modules for most Internet protocols, including the following:

- Web Distributed Authoring and Versioning (WebDAV)
- Internet Message Access Protocol version 4rev1 (IMAP4)
- Lightweight Directory Access Protocol (LDAP)
- OLE DB
- Post Office Protocol version 3 (POP3)
- SMTP

For more information about the ESP tool or to download ESP, see the Microsoft Exchange Web site at <http://go.microsoft.com/fwlink/?LinkId=1709>.

### Monitoring Performance

When you deploy active/active or active/passive Exchange 2000 Server clusters, it is important to proactively monitor the clusters. Use System Monitor and Event Viewer to view the overall performance of your server and the performance of Exchange 2000 Server.

The following counters are virtual memory counters that are especially important to monitor when deploying Exchange 2000 Server clusters:

- **MSExchangeIS\VM Largest Block Size** Displays the size in bytes of the largest free block of virtual memory. This counter is a line that slopes down as virtual memory is consumed. When this counter drops below 32 MB, Exchange 2000 logs a warning in the event log (Event ID=9582) and logs an error if this drops below 16 MB. It is important to monitor this counter to ensure that it stays above 32 MB.
- **MSExchangeIS\VM Total 16MB Free Blocks** Displays the total number of free virtual memory blocks that are greater than or equal to 16 MB. This line forms a pyramid as you monitor it. It starts with one block of virtual memory greater than 16 MB and progresses to smaller blocks greater than 16 MB. Monitoring the trend on this counter should allow a system administrator to predict when the number of 16 MB blocks is likely to drop below 3, at which point restarting all the services on the node is recommended.
- **MSExchangeIS\VM Total Free Blocks** Displays the total number of free virtual memory blocks regardless of size. This line forms a pyramid as you monitor it. This

counter can be used to measure the degree to which available virtual memory is being fragmented. The average block size is the Process\Virtual Bytes\STORE instance divided by MExchangeIS\VM Total Free Blocks.

- **MExchangeIS\VM Total Large Free Block Bytes** Displays the sum in bytes of all the free virtual memory blocks that are greater than or equal to 16 MB. This line slopes down as memory is consumed.

When you monitor these counters, the most important thing to monitor is that **VM Total Large Free Block Bytes** always exceeds 32 MB. If a node in the cluster drops below 32 MB, fail over the Exchange Virtual Servers, restart all of the services on the node, and then fail back the Exchange Virtual Servers.

The Microsoft Exchange Information Store service logs the following events if the virtual memory for your Exchange 2000 Server becomes excessively fragmented:

```
EventID=9582
```

```
Severity=Warning
```

```
Facility=Perfmon
```

```
Language=English
```

```
The virtual memory necessary to run your Exchange server is fragmented in such a way that performance may be affected. It is highly recommended that you restart all Exchange services to correct this issue.
```

**Note** This warning is logged if the largest free block is smaller than 32 MB.

```
EventID=9582
```

```
Severity=Error
```

```
Facility=Perfmon
```

```
Language=English
```

```
The virtual memory necessary to run your Exchange server is fragmented in such a way that normal operation may begin to fail. It is highly recommended that you restart all Exchange services to correct this issue.
```

**Note** This error is logged if the largest free block is smaller than 16 MB.

For more information about System Monitor and Event Viewer, see Windows 2000 Help.

## Exchange 2000 Server Cluster Failover Performance

Many factors determine the speed of Exchange 2000 Server cluster failovers. If an Exchange Virtual Server must fail over, certain tasks must be accomplished to complete the failover. Understanding these tasks helps you configure your Exchange 2000 Server clusters for the fastest possible failover.

## Extensible Storage Engine Log Checkpoint Depth

Storage group databases write new transactions to a log and then update the database when it is efficient to do so. The maximum number of logs that can be written before the data is committed to the database is called the log checkpoint depth. By default, this depth is 20 MB. It is possible for the Microsoft Exchange Information Store service to write 20 MB of logs before it writes that information into the actual database.

**Note** When an Exchange Virtual Server moves or fails over from one node to another node, all transactions are committed to the database before the Information Store resource goes offline.

In general, reducing the checkpoint depth decreases the amount of data that must be committed to the database from the log when shutdown is initiated.

You can modify the Extensible Storage Engine checkpoint depth using ADSI Edit. You modify the checkpoint depth by modifying the Extensible Storage Engine checkpoint depth attribute on a per storage group basis.

### To modify the Extensible Storage Engine checkpoint depth

1. Run ADSI Edit. ADSI Edit is provided with the Windows 2000 Server Support Tools.
2. In the console tree, expand **Configuration Container, CN=Configuration, CN=Services, CN=Microsoft Exchange, CN=<your organization name>, CN=Administrative Groups, CN=<your administrative group>, CN=Servers, CN=<server name>, CN=Information Store**, and **CN=<your storage group>**.
3. Right-click **CN=<your storage group>**, and then click **Properties**.
4. In **Select a property to view**, select **msExchESEParamCheckpointDepthMax**.
5. In **Edit Attribute**, type the size you want to set the maximum cache size to (in 5-MB increments entered in bytes), and then click **Set**. For example, a 5-MB checkpoint depth is 5,242,880 bytes, and a 10-MB checkpoint depth is 10,485,760 bytes. The minimum size you can enter is 5 MB. If you enter a value over 20 MB, Exchange Information Store resource offline times increase.
6. Quit ADSI Edit, and then restart the Microsoft Exchange Information Store service.

**Note** In a multiple domain controller environment, this attribute change requires time to replicate for Exchange 2000 Server to reflect the change. In addition, decreasing the checkpoint depth from 20 MB to 5 MB causes a performance penalty of 5 percent. This is because the reduced checkpoint depth forces the Extensible Storage Engine to commit data from the logs more often, which causes disk and CPU performance penalties.

### Microsoft Exchange Information Store Service Connections

The number of connections into the Microsoft Exchange Information Store service affects failover time. The Microsoft Exchange Information Store service performs cleanup routines before it releases and allows failover to occur. An unloaded server that takes 100 seconds to fail over fully takes 120 seconds to fail over with 3,000 simultaneous Microsoft Outlook® Web Access or Microsoft Outlook connections.

## SMTP Queue Size

The number of messages in the SMTP queue is also a factor in the time it takes an Exchange Virtual Server to fail over from one cluster node to another. This time can be significant if the queue size is over 1,000 messages. To reduce this time, modify the **Max Handle Threshold** registry key.

For more information, see the "Tuning Exchange 2000 Server Clusters" section later in this document.

## Evenly Distributing Threads across SMTP, IMAP4, and POP3

Five work queues in SMTP use a pool of threads known as the ATQ threads. By design, two of these queues can use up to 90 percent of the available threads. SMTP shares this thread pool with the process that accepts POP3 and IMAP4 requests. Therefore, in a moderate to high load scenario, a situation can arise where the SMTP service keeps resources away from POP3 and IMAP4 services.

When this situation occurs, the IMAP4 and POP3 resources may fail their "Is Alive" poll from Cluster Service, and may eventually cause the Exchange Virtual Server to fail over to another node.

You can configure Exchange 2000 Server to reserve adequate threads for IMAP4 and POP3 by limiting the percentage of threads that the SMTP service can use. To configure this limit, increase the overall number of threads available to IIS.

For more information about distributing threads across SMTP, IMAP3, and POP3, see the "Tuning Exchange 2000 Server Clusters" section later in this document.

**Note** This change limits the threads used by the SMTP service and, as a result, decreases SMTP performance when processing large queues. An increase in the number of overall threads available to IIS also causes an increase in memory usage.

## Tuning Exchange 2000 Server Clusters

Use the following registry changes to help you tune the performance of your computer running Exchange 2000.

**Important** This section contains information about editing the registry. Before you edit the registry, ensure you understand how to restore it if a problem occurs. For information on restoring the registry, see the "Restore the Registry" Help topic in Regedit.exe or Regedt32.exe.

### SMTP Percentage of Threads and Additional Threads per Processor

To reserve adequate threads for POP3 and IMAP4, limit the percentage of threads the SMTP service is allowed to use. To make this change, you have to increase the overall number of threads available to IIS.

Two registry values allow you to control the SMTP percentage of threads and additional threads per processor (see Table 7).

**Warning** Using Registry Editor incorrectly can cause serious problems that may require you to reinstall your operating system. There is no guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk. For information about how to edit the registry, see the "Change

Keys and Values" Help topic in Registry Editor (Regedit.exe) or the "Add and Delete Information in the Registry" and "Edit Registry Information" Help topics in Regedt32.exe. Note that you should back up the registry before you edit it. If you are running Microsoft® Windows NT® or Microsoft Windows 2000, you should also update your Emergency Repair Disk (ERD).

**Table 7** This table shows SMTP percentage of threads and additional threads per processor registry keys

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SMTPSVC\Queuing.
Parameter	MaxPercentPoolThreads (REG_DWORD).
Description	Maximum percentage of ATQ threads that each queue will request.
Default setting	Not present, but will default to 0x5A (90).
When to change	Adjust when high SMTP activity causes the POP3 and IMAP4 resources to fail in Cluster Administrator.
Recommended setting	Use the following formula to determine the settings:  $\text{MaxPercentPoolThreads} = 90 \div (2 \times \text{Number of SMTP Virtual Server Instances across all Exchange Virtual Servers that can run on this node}).$ <p>For example, assuming a cluster with two SMTP Virtual Server Instances = <math>90 \div (2 \times 2) = 22</math>.</p> <p><b>Note</b> If you get a decimal point in the calculation, round down to the next integer.</p>
Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SMTPSVC\Queuing.
Parameter	AdditionalPoolThreadsPerProc (REG_DWORD).
Description	Additional pool per processor threads when SMTP is started.
Default setting	Not present, but will default to 0x6 (6).
When to change	Adjust when high SMTP activity causes the POP3 and IMAP4 instances to fail in Cluster Administrator. Change in conjunction with MaxPercentPoolThreads.

Recommended setting Use the following formula to determine the settings:

$$\text{AdditionalPoolThreadsPerProc} = ((9 \div (\text{MaxPercentPoolThreads} \div 100)) - 4) \div 2$$

For example, assuming a cluster with two SMTP Virtual Server Instances  $((9 \div (0.22)) - 4) \div 2 = (41 - 4) \div 2 = 18$

**Note** If you get a decimal point in the calculation, round down to the next integer. If you get a number greater than 200, assign that number to the following registry value:

HKEY\_LOCAL\_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters\PoolThreadLimit

### Maximum Handle Threshold

The default value for this setting in Exchange 2000 Server with SP3 is 1,000 handles per 256 MB of memory with up to 10,000 handles total. Setting this value below 10,000 decreases SMTP resource failover times in large queue situations. This registry key forces the SMTP service to keep fewer messages "open," and as a result, when a failover is required, fewer messages in the queue will need to be cleared from memory and written back to disk.

For the steps necessary to modify the Maximum Handle Threshold registry key, see the Microsoft Knowledge Base article Q271084, "XGEN: Exchange 2000 Server SMTP Optimized with Maximum Handle Threshold Registry Key."

**Note** Decreasing the Maximum Handle Threshold value causes the SMTP service to use a less efficient method for processing the queues, which increases both CPU and disk usage in large queue situations.

## Disaster Recovery

Clustering provides a mechanism for moving resources between cluster nodes when a disaster occurs. When a single server fails, clustering moves Exchange 2000 resources to another server in the cluster so that services remain available to users. After a disaster, you should attempt to repair the damaged server. If that does not work, you may need to replace a damaged cluster node, restore or rebuild a cluster node from backups, restore a shared disk resource from backups, or recover the entire cluster. Although the disaster recovery processes for restoring Exchange 2000 clusters are similar to the processes for restoring data on stand-alone Exchange servers, there are important differences.

**Important** For detailed conceptual information and step-by-step procedures about backing up and restoring Exchange 2000 clusters, see the "Backing Up Exchange 2000 Clusters" and "Restoring Exchange 2000 Clusters" sections of the technical paper *Disaster Recovery for Microsoft Exchange 2000 Server* (<http://go.microsoft.com/fwlink/?LinkId=1714>), before proceeding with any cluster backup and recovery steps.

## Identifying the Cause of a Failure

An important task in disaster recovery processes for Exchange 2000 clusters is identifying what caused a particular resource to fail. When a failure occurs in an Exchange cluster, you should first determine if the failure is on a single node (which indicates that there are problems with the node's files) or on every node (which indicates that there are problems with the cluster's objects or the shared cluster resources). To determine the cause of the failure, search the event logs within Event Viewer. You can also search for resolutions in the Microsoft Knowledge Base at <http://support.microsoft.com/>.

If you are still unable to determine the cause of the failure, you can perform the repair options listed in "Repairing Windows 2000" or "Repairing Exchange 2000" in the technical article *Disaster Recovery for Exchange 2000 Server* available at <http://go.microsoft.com/fwlink/?LinkId=1714>. If repairing the node or entire cluster is unsuccessful, you must consider replacing the node or recovering the node, cluster, or resources (such as the quorum disk resource or Exchange mailbox and public folder stores).

## Backing Up Data on an Exchange 2000 Server Cluster Node

Securing the data on your Exchange 2000 clusters requires establishing a proper and thorough backup plan. To back up the important data on the nodes of your Exchange 2000 clusters, you can use Windows 2000 Backup. You can also use third-party backup solutions to meet your backup needs. For information about third-party backup solutions, see "Exchange 2000 Third-party Solutions" on the Microsoft Exchange Web site at <http://go.microsoft.com/fwlink/?LinkId=5225>.

To ensure the data on your Exchange 2000 clusters is secure, it is also important to maintain informational records about your clusters.

To secure the data in your clusters, you must do the following:

- Back up Windows 2000 in each cluster node
- Back up the quorum disk resource of each cluster
- Back up all of the Exchange databases on your shared disk resources
- Maintain informational records about your cluster configuration

For information about backing up this data and maintaining records about your cluster configuration, see the technical article *Disaster Recovery for Microsoft Exchange 2000* available at <http://go.microsoft.com/fwlink/?LinkId=1714>.

## Recovering an Exchange 2000 Server Cluster

Recovering from disasters that affect the nodes of your Exchange 2000 clusters can be as simple as replacing a node with a stand-by recovery server, or it can be as difficult as rebuilding an entire cluster from the beginning. If you have a proper and thorough backup plan in place, you can recover from almost any disaster that affects your Exchange organization.

You may need to do the following to recover from disasters that affect your Exchange 2000 clusters:

- Replace damaged cluster nodes
- Restore or rebuild a cluster node from backups
- Restore shared disk resources
- Restore quorum disk resource
- Restore Exchange databases
- Recover an entire Exchange 2000 cluster

For information about recovering from disasters that affect your Exchange 2000 clusters, see the technical article *Disaster Recovery for Microsoft Exchange 2000* at <http://go.microsoft.com/fwlink/?LinkId=1714>.

## Additional Resources

The following resources have additional information that will help you in your setup and configuration of clusters.

### Technical Articles and Tools

- *Introducing Windows 2000 Clustering Technologies*  
<http://go.microsoft.com/fwlink/?LinkId=1719>
- *Exploring Windows 2000 Clustering Technologies*  
<http://go.microsoft.com/fwlink/?LinkId=1720>
- *Step-by-Step Guide to Installing Cluster Service*  
<http://go.microsoft.com/fwlink/?LinkId=266>
- *Windows 2000 Clustering: Performing a Rolling Upgrade*  
<http://go.microsoft.com/fwlink/?LinkId=1718>
- *Storage Solutions for Exchange 2000*  
<http://go.microsoft.com/fwlink/?LinkId=1715>
- *Exchange 2000 Front-End and Back-End Topology*  
<http://go.microsoft.com/fwlink/?LinkId=4721>
- *Microsoft Exchange 2000 Front-End Server and SMTP Gateway Hardware Scalability Guide*  
<http://go.microsoft.com/fwlink/?LinkId=1713>
- *Microsoft Exchange 2000 Server Back-End Mailbox Scalability*  
<http://go.microsoft.com/fwlink/?LinkId=1711>
- *Exchange 2000 Server LoadSim Tool*  
<http://go.microsoft.com/fwlink/?LinkId=1710>
- *Exchange 2000 Server ESP Tool*  
<http://go.microsoft.com/fwlink/?LinkId=1709>
- *Exchange 2000 Server Capacity and Topology Calculator*  
<http://go.microsoft.com/fwlink/?LinkId=1716>
- *Disaster Recovery for Exchange 2000 Server*  
<http://go.microsoft.com/fwlink/?LinkId=1714>

## Microsoft Knowledge Base Articles

There are many Microsoft Knowledge Base articles available on Exchange 2000 clusters. Be sure to consult these articles if you experience problems with your Exchange 2000 Server clusters.

### To search for Knowledge Base articles related to Exchange 2000 Server clusters:

1. In your browser, go to <http://search.support.microsoft.com/> default.
2. In the **Search the Knowledge Base** drop-down box, select **Exchange 2000 Server**.
3. In **For solutions containing (optional)**, enter your keywords.
4. Click **Search now**.

The following specific Microsoft Exchange Knowledge Base articles relate to the topics of this technical article and provide more information about deploying and maintaining your Exchange 2000 clusters.

### Setup-related

[Q262356](#) - XADM: Options Not Supported During an Unattended Installation

[Q270668](#) - XADM: Exchange 2000 Setup Fails with 0XC103798A

[Q277908](#) - XWEB: How to Enable the Change Password Option in Exchange 2000 OWA Clustering

[Q278244](#) - XADM: Error Message Occurs When Configuring a New System Attendant Resource on a Cluster

[Q285137](#) - XADM: "C0072030" Error Message Occurs When You Create a System Attendant Resource on a Cluster Server

[Q305145](#) - HOW TO: Remove the IFS Mapping for Drive M in Exchange 2000 Server

[Q316886](#) - HOW TO: Migrate from Exchange Server 5.5 to Exchange 2000 Server

[Q320007](#) - XADM: Permissions That Were Modified Manually Are Reset to the Default Values

### Maintenance-related

[Q271407](#) - XADM: Default SMTP Log File Directory Is Incorrect for Clustered Servers

[Q254288](#) - XADM: When You Create the System Attendant on a Cluster, an RPC Error Occurs

[Q263060](#) - XADM: Full Exchange Admin Cannot Create a New Store on the Second Virtual Server in a Cluster

[Q264624](#) - XADM: Content Indexing Logs Catastrophic Error When Shutting Down

[Q266367](#) - XADM: Extensible Storage Engine 98 Error Codes -1051 to -999999

[Q266689](#) - XADM: The "ESEUTIL /CC" Command Does Not Work on Cluster Server

[Q271651](#) - XADM: How to Add a Node to an Exchange 2000 Cluster

[Q288598](#) - XADM: Event ID 9175 Logged by System Attendant Consistently on a Cluster

[Q293510](#) - XADM: How to Add an Exchange 2000 Virtual Server to a Cluster Server

**Performance-related**

[Q296073](#) - XADM: Monitoring for Exchange 2000 Memory Fragmentation

**Remove Exchange-related**

[Q224975](#) - XADM: Differences Between Move Mailbox Methods

[Q262895](#) - XADM: Removing Exchange 2000 Does Not Remove "Organization"

[Q288150](#) - XADM: How to Rehome Public Folders in Exchange 2000

[Q297393](#) - HOW TO: Programmatically Move an Exchange 2000 Mailbox Using CDOEXM in Visual C++

[Q309511](#) - XADM: Exchange Management Service Remains After You Remove Exchange 2000 Server

[Q309751](#) - XADM: You Cannot Delete Files on a Shared Disk If Exchange Cluster Resources Have Been Deleted

**For more information:** <http://www.microsoft.com/exchange/default.asp>

Does this paper help you? Please give us your feedback. On a scale of 1 (poor) to 5 (excellent), how do you rate this paper?

<mailto:exchdocs@microsoft.com?subject=Feedback: Deploying Microsoft Exchange 2000 Server Clusters>