

Dell™ PowerEdge™ Cluster
FE500W Systems

Installation and Troubleshooting Guide

Notes, Notices, and Cautions



NOTE: A NOTE indicates important information that helps you make better use of your computer.



NOTICE: A NOTICE indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.



CAUTION: A CAUTION indicates a potential for property damage, personal injury, or death.

Information in this document is subject to change without notice.

© 2005–2007 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, *PowerVault*, and *OpenManage* are trademarks of Dell Inc.; *Microsoft*, *Windows*, *Windows Server*, *Active Directory*, and *Windows NT* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries; *EMC* is a registered trademark and *Navisphere*, *Navisphere Agent*, *Navisphere Manager*, *Access Logix*, *ControlCenter*, *MirrorView*, *SAN Copy*, and *SnapView* are trademarks of EMC Corporation; *Intel* and *Pentium* are registered trademarks of Intel Corporation.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

December 2007

Rev. A01

Contents

1	Introduction	9
	Overview	9
	Virtual Servers and Resource Groups	9
	Quorum Resource	10
	Cluster Solution	10
	Operating System	11
	Shared Storage Systems	11
	Dell EMC Storage Management Software	13
	Fibre Channel Overview	15
	Supported Cluster Configurations	17
	Direct-Attached Cluster	17
	SAN-Attached Cluster	18
	System Requirements	18
	Cluster Nodes	19
	Cluster Storage	20
	Other Documents You May Need	21
2	Cabling Your Cluster Hardware	23
	Cabling the Mouse, Keyboard, and Monitor	23
	Cabling the Power Supplies	23

Cabling Your Cluster For Public and Private Networks	26
Cabling the Public Network.	27
Cabling the Private Network	28
NIC Teaming.	29
Fibre Channel Cable Connectors	29
Cabling the Storage Systems	31
Cabling Storage for Your Direct-Attached Cluster.	31
Cabling Storage for Your SAN-Attached Cluster.	35
3 Preparing Your Systems for Clustering	47
Before You Begin.	47
Installation Overview	47
Selecting a Domain Model	49
Configuring the Nodes as Domain Controllers	49
Configuring Internal Drives in the Cluster Nodes.	50
Installing and Configuring the Microsoft Windows Operating System.	50
Configuring Windows Networking	52
Assigning Static IP Addresses to Cluster Resources and Components	52
Configuring IP Addresses for the Private Network.	53
Verifying Communications Between Nodes	56
Configuring the Internet Connection Firewall	56

Installing the Fibre Channel HBAs	57
Installing the Fibre Channel HBA Drivers	57
Implementing Zoning on a Fibre Channel Switched Fabric	57
Using Zoning in SAN Configurations Containing Multiple Hosts	58
Using Worldwide Port Name Zoning	58
Installing and Configuring the Shared Storage System	60
Access Logix	61
Access Control	62
Storage Groups	63
Navisphere Manager	64
Navisphere Agent	65
EMC PowerPath	65
Enabling Access Logix and Creating Storage Groups Using Navisphere 6.x.	66
Configuring the Hard Drives on the Shared Storage System(s)	67
Updating a Dell EMC Storage System for Clustering	74
Installing and Configuring MSCS	74
Installing and Configuring Microsoft Cluster Service (MSCS) with Windows 2000	74
Configuring Microsoft Cluster Service (MSCS) with Windows Server 2003	75
Verifying Cluster Readiness	77
Installing Applications in the Cluster Group	78
Installing the Quorum Resource	78
Creating a LUN for the Quorum Resource	78

	Configuring Windows 2000 Cluster Networks	79
	Configuring Cluster Networks Running Windows Server 2003.	79
	Verifying MSCS Operation	80
	Verifying Cluster Functionality	80
	Verifying Cluster Resource Availability	80
	Troubleshooting Failed Resources	81
4	Installing Your Cluster Management Software	83
	Microsoft Cluster Administrator	83
	Launching Cluster Administrator on a Cluster Node	83
	Running Cluster Administrator on a Remote Console.	83
	Launching Cluster Administrator on a Remote Console.	84
	Installing Cluster Administrator for Windows Clusters on a Remote Console	84
5	Using MSCS	87
	Cluster Objects	87
	Cluster Networks	87
	Preventing Network Failure.	87
	Node-to-Node Communication	88
	Network Interfaces	88

Cluster Nodes	88
Forming a New Cluster	89
Joining an Existing Cluster	89
Cluster Resources	89
Setting Resource Properties	90
Resource Dependencies	90
Setting Advanced Resource Properties	91
Resource Parameters	91
Quorum Resource	92
Resource Failure	93
Resource Dependencies	94
Creating a New Resource	94
Deleting a Resource	95
File Share Resource Type	96
Configuring Active and Passive Cluster Nodes	96
Failover Policies	99
Windows 2000 Advanced Server Cluster Configurations	99
Windows Server 2003 Cluster Configurations	99
Failover and Failback Capabilities	105
6 Upgrading to a Cluster Configuration	107
Before You Begin	107
Supported Cluster Configurations	107
Completing the Upgrade	107
Upgrading Your Operating System	108
Performing a Rolling Upgrade	108

7	Maintaining Your Cluster	113
	Adding a Network Adapter to a Cluster Node	113
	Changing the IP Address of a Cluster Node on the Same IP Subnet	114
	Uninstalling MSCS From Clusters Running Microsoft Windows 2000 Advanced Server	115
	Removing Nodes From Clusters Running Microsoft Windows Server 2003	115
	Running chkdsk /f on a Quorum Resource	116
	Recovering From a Corrupt Quorum Disk	116
	Changing the MSCS Account Password in Windows Server 2003	117
	Reformatting a Cluster Disk	118
A	Troubleshooting	121
B	Cluster Data Form	129
C	Zoning Configuration Form	131

Introduction

This document provides information for installing and managing your cluster solution. It is intended for experienced IT professionals who need to configure the cluster solution, and for trained service technicians who perform upgrade and maintenance procedures. This document also addresses readers who are new to clustering.

Overview

Clustering uses specific hardware and software to join multiple systems together to function as a single system and provide an automatic failover solution. If one of the clustered systems (also known as cluster nodes, or nodes) fails, resources running on the failed system are moved (or failed over) to one or more systems in the cluster by the Microsoft® Cluster Service (MSCS) software. MSCS is the failover software component in specific versions of the Windows® operating system.

When the failed system is repaired and brought back online, resources automatically transfer back (or fail back) to the repaired system or remain on the failover system, depending on how MSCS is configured. For more information, see "Configuring Active and Passive Cluster Nodes" on page 96.



NOTE: Reference to Windows Server™ 2003 in this guide implies reference to both Windows Server 2003, Enterprise and Windows Server 2003 Enterprise x64 Editions, unless explicitly stated.

Virtual Servers and Resource Groups

In a cluster environment, users do not access a physical server; they access a virtual server, which is managed by MSCS. Each virtual server has its own IP address, name, and hard drive(s) in the shared storage system. MSCS manages the virtual server as a *resource group*, which contains the cluster resources. Ownership of virtual servers and resource groups is transparent to users. For more information on resource groups, see "Cluster Resources" on page 89.

When MSCS detects a failed server node or failed application, MSCS moves the failed resource group(s) to one or more server nodes and remaps the virtual server(s) to the new network connection(s). Users of an application in the virtual server experience only a momentary delay in accessing resources while MSCS re-establishes a network connection to the virtual server and restarts the application.

Quorum Resource

A single shared disk, which is designated as the quorum resource, maintains the configuration data (including all the changes that have been applied to a cluster database) necessary for recovery when a node fails.

The quorum resource can be any resource with the following attributes:

- Enables a single node to gain and defend its physical control of the quorum resource
- Provides physical storage that is accessible by any node in the cluster
- Uses the Microsoft Windows NT[®] file system (NTFS)

See "Quorum Resource" on page 92 and the MSCS online documentation for more information.



NOTE: PowerEdge™ FE clusters do not support the Majority Node Set Quorum resource type.

Cluster Solution

Your cluster implements two-node to eight-node clustering and provides the following features:

- 2 Gb/sec Fibre Channel technology
- High availability of resources to network clients
- Redundant paths to the shared storage
- Failure recovery for applications and services
- Flexible maintenance capabilities, allowing you to repair, maintain, or upgrade a node or storage system without taking the entire cluster offline

Operating System

Table 1-1 provides an overview of the supported operating systems. See your operating system documentation for a complete list of features.





 **NOTE:** Some of the core services are common to all the operating systems.

Table 1-1. Windows Operating System Features

Windows 2000 Advanced Server	Windows Server 2003, Enterprise Edition	Windows Server 2003, Enterprise x64 Edition
Supports two-node clusters	Supports up to eight nodes per cluster	Supports up to eight nodes per cluster
Supports up to 8 GB of RAM per node	Supports up to 32 GB of RAM per node	Supports up to 1 TB RAM per node
Cluster configuration and management using Control Panel utilities	Cluster configuration and management using Configure Your Server (CYS) and Manage Your Server (MYS) wizards Metadirectory Services	Cluster configuration and management using Configure Your Server (CYS) and Manage Your Server (MYS) wizards Metadirectory Services

 **NOTE:** The amount of RAM supported per node also depends on your cluster platform.


 **NOTE:** Running different operating systems in a cluster is supported only during a rolling upgrade. You cannot upgrade to Windows Server 2003, Enterprise x64 Edition. Only a new installation is permitted for Windows Server 2003, Enterprise x64 Edition.

 **NOTE:** MSCS and Network Load Balancing (NLB) features cannot coexist on the same node, but can be used together in a multitiered cluster. For more information, see the Dell PowerEdge Clusters website at www.dell.com/clusters or the Microsoft website at www.microsoft.com.

Shared Storage Systems

Cluster nodes can share access to external storage systems; however, only one of the nodes can own any RAID volume in the external storage system at any time. MSCS controls which node has access to each RAID volume in the shared storage system.

Your cluster supports Dell | EMC® CX-Series storage systems in a direct-attached or SAN-attached environment.

 **NOTE:** EMC® Access Logix™ software is required when the storage system is connected to two or more clusters, two or more non-clustered systems, or a combination of both clustered and non-clustered systems. Access Logix software is not required for stand-alone systems or single-cluster configurations.

Each storage system in the cluster is centrally managed by one host system (also called a *management station*) running EMC® ControlCenter™ Navisphere Manager™—a centralized storage management application used to configure Dell | EMC storage systems. Using a GUI, you can select a specific view of your storage arrays, as shown in Table 1-2.

Table 1-2. Navisphere Manager Storage Views

View	Description
Storage	Shows the logical storage components and their relationships to each other and identifies hardware faults.
Hosts	Shows the host system's storage group and attached logical unit numbers (LUNs).
Monitors	Shows all Event Monitor configurations, including centralized and distributed monitoring configurations.

You can use Navisphere Manager to perform tasks such as creating RAID arrays, binding LUNs, and downloading firmware.

Optional software for the shared storage systems include:

- EMC MirrorView™ — Provides synchronous or asynchronous mirroring between two storage systems.
- EMC SnapView™ — Captures point-in-time images of a LUN for backups or testing without affecting the contents of the source LUN.
- EMC SAN Copy™ — Moves data between Dell | EMC storage systems without using host CPU cycles or LAN bandwidth.

See "Installing and Configuring the Shared Storage System" on page 60 for more information about Navisphere Manager, Access Logix, MirrorView, SnapView, and SAN Copy™.

Dell | EMC Storage Management Software

The storage systems work together with the following hardware components:

- Processor enclosure (DPE or SPE) - Configured with storage processors that control the RAID arrays in the storage system and provide storage functionalities such as snapshots, LUN masking, and remote mirroring. Supported processor enclosures vary in performance and configuration.
- Disk array enclosure (DAE2) - Provides additional storage and is attached to the processor enclosure.
- Standby power supply (SPS) - Provides backup power to protect the integrity of the storage processor write cache. The SPS is connected to the storage processor and the disk enclosure that has the core software preinstalled.

Cluster nodes can share access to the storage systems; however, only one of the nodes can own any RAID volume in the external storage system at any time. MSCS controls which node has access to each RAID volume in the shared storage system.

This section provides information about the software your cluster uses to manage the communications between the nodes and the storage systems in a SAN environment.

EMC PowerPath

PowerPath™ detects and re-establishes a failed connection to a processor enclosure by automatically rerouting I/Os through an alternate path. PowerPath also provides load balancing of data from the server to the storage system.

Navisphere Manager

Navisphere Manager™ provides centralized storage management and configuration, allowing you to configure and manage the disks and components in one or more shared storage systems. Navisphere Manager is installed on the storage systems or a host system that has network access to the storage system.

Navisphere Agent

Navisphere Agent™ provides an interface between the host system and the storage system, allowing Navisphere Manager to send and receive information to and from the storage system connected to a host system.

Access Logix (optional)

Access Logix™ enables multiple nodes and servers to share a storage system. It restricts server access to specific volumes on a shared storage system and protects data from unauthorized access.

Access Logix is required when your PowerEdge system contains server modules that are configured in heterogeneous configurations. These configurations include:

- Two or more stand-alone systems
- Two or more clusters
- Any combination of server modules configured as cluster nodes and stand-alone systems.

MirrorView (optional)

MirrorView™ automatically duplicates primary storage system data from a cluster or stand-alone system to a secondary storage system. It can be used in conjunction with SnapView and is managed from within Navisphere Manager.

SnapView (optional)

SnapView™ captures images of a LUN and retains the images independently of subsequent changes to the files. The images can be used to share LUNs with another system without affecting the contents of the source LUN.

SnapView creates copies of LUNs using either snapshots or clones. Snapshots are virtual copies that create an image of the source LUN at the time the snapshot was created. This snapshot is retained independently of subsequent changes to the source LUN. Clones are duplicate copies of a source LUN. You can use snapshots and clones to facilitate backups or to allow multiple hosts to access data without affecting the contents of the source LUN.



NOTE: Each snapshot or clone must be accessed from a different host.

SnapView, which is installed on the storage processors as a nondisruptive upgrade, can be used in conjunction with MirrorView and is managed from within Navisphere Manager.

SAN Copy (optional)

SAN Copy allows you to move data between storage systems without using host processor cycles or LAN bandwidth. It can be used in conjunction with SnapView or MirrorView and is managed from within Navisphere Manager.

Fibre Channel Overview

This section provides a brief overview of Fibre Channel and the elements of a Fibre Channel network.

Fibre Channel Protocol

Fibre Channel provides high-speed data transfer between the nodes and storage systems, allows multiple server systems to share one or more storage systems, and provides long-distance connectivity and high bandwidth. Fibre Channel switches can be linked together using interswitch links (ISLs). These ISLs use two Fibre Channel ports to connect the switches together. By employing long-wave fiber optic cable between cascaded switches, systems up to 10 km from the shared storage array can access data as if they are directly attached.

Implementing Fibre Channel technology in a cluster provides the following advantages:

- **Flexibility** — Fibre Channel allows a distance of up to 10 km between switches without degrading the signal.
- **Availability** — Fibre Channel components use redundant connections, providing multiple data paths and greater availability for clients.
- **Connectivity** — Fibre Channel allows more device connections than SCSI. Because Fibre Channel devices are hot-pluggable, you can add or remove devices from the nodes without bringing down the cluster.

Fibre Channel Switch Fabric

A Fibre Channel switch fabric consists of one or more switches that provide a connection from one device (sender) to another device or switch (receiver) on the network. If the data is sent to another switch, the process repeats until a connection is established between the sender and the receiver.

SAN

In addition to direct-attached connections, nodes can also connect to storage systems through a SAN—a high-performance network storage solution. A SAN bypasses traditional network bottlenecks, providing a high-speed, highly available data consolidation solution.



NOTE: Your SAN may require additional hardware and software components that are not listed in this document or the *Platform Guide*. See the EMC Support Matrix located at www.emc.com for information about SAN-compliant hardware and software components.

Zones

Fibre Channel fabrics enable you to set up barriers between different devices and operating environments. These barriers create logical fabric subsets, called *zones*, which divide a fabric into smaller groups or components, regardless of their proximity to one another.

By implementing switch zoning together with Access Logix™, you can attach multiple clusters or a combination of clusters and stand-alone systems to a storage system in a SAN.



NOTE: Your cluster supports only single-initiator zoning for connecting clusters to the storage systems in a switched environment. Each physical Host Bus Adapter (HBA) port has a unique worldwide name (WWN), and is treated as an initiator. A separate zone is created for each initiator. This zone includes the HBA port, one or more Storage Processor (SP) ports on the storage system(s), and may include a tape library.

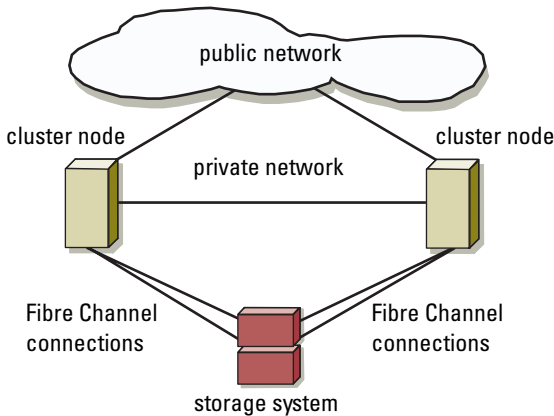
Supported Cluster Configurations

Direct-Attached Cluster

In a direct-attached cluster, both nodes of the cluster are directly attached to a single storage system. In this configuration, the RAID controllers (or storage processors) on the storage systems are connected by cables directly to the Fibre Channel HBA ports in the nodes.

Figure 1-1 shows a basic direct-attached, single-cluster configuration.

Figure 1-1. Direct-Attached, Single-Cluster Configuration



EMC PowerPath Limitations in a Direct-Attached Cluster


EMC PowerPath provides failover capabilities and multiple path detection as well as dynamic load balancing between multiple ports on the same storage processor. However, direct-attached clusters supported by Dell connect to a single port on each storage processor in the storage system. Because of the single port limitation, PowerPath can provide only failover protection, not load balancing, in a direct-attached configuration.

SAN-Attached Cluster

In a SAN-attached cluster, all of the nodes are attached to a single storage system or to multiple storage systems through a SAN using redundant switch fabrics. SAN-attached clusters are superior to direct-attached clusters in configuration flexibility, expandability, and performance.

Figure 1-2 shows a SAN-attached cluster.

Figure 1-2. SAN-Attached Cluster



System Requirements

Your cluster requires the following components:

- Servers (nodes)
- Storage
- Interconnects (private network)
- Client network connections (public network)
- Operating system and storage management software

Cluster Nodes

Table 1-3 lists the hardware requirements for the cluster nodes.

Table 1-3. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	Two supported Dell PowerEdge systems running the Microsoft® Windows® 2000, Advanced Server operating system. OR Two to eight PowerEdge systems running the Windows Server 2003 operating system.
RAM	At least 256 MB of RAM installed on each cluster node for Windows 2000 Advanced Server and Windows Server 2003, Enterprise Edition. At least 512 MB of RAM installed on each cluster node for Windows Server 2003, Enterprise x64 Edition.
HBA ports	Two Fibre Channel HBAs per node, unless the server includes an integrated dual-port Fibre Channel HBA. Where possible, place the HBAs on separate PCI buses to improve availability and performance. See the <i>Platform Guide</i> for information about supported systems, HBAs, and PCI slot configuration guidelines.
NICs	At least two NICs: one NIC for the public network and another NIC for the private network. NOTE: It is recommended that the NICs on each public network are identical, and that the NICs on each private network are identical.
Internal disk controller	One controller connected to at least two internal hard drives for each node. Use any supported RAID controller or SCSI adapter. Two hard drives are required for mirroring (RAID 1) and at least three are required for disk striping with parity (RAID 5). NOTE: It is strongly recommended that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives.

Cluster Storage

Table 1-4 lists supported storage systems and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems. Table 1-5 lists hardware requirements for the disk processor enclosures (DPE), storage processor enclosures (SPE), disk array enclosures (DAE2), and standby power supplies (SPS).

Table 1-4. Cluster Storage Requirements

Hardware Components	Requirement
Supported storage systems	One to four supported Dell EMC storage systems. See Table 1-5 for specific storage system requirements.
Cluster nodes	All nodes must be directly attached to a single storage system or attached to one or more storage systems through a SAN.
Multiple clusters and stand-alone systems	Can share one or more supported storage systems using optional software that is available for your storage system. See "Dell EMC Storage Management Software" on page 13.

Table 1-5. Dell | EMC Storage System Requirements

Processor Enclosure	Minimum Required Storage	Possible Storage Expansion	SPS
CX300 DPE	At least five and up to 15 internal hard drives	Up to three DAE2 with a maximum of 15 hard drives each	Two per DPE
CX500 DPE	At least five and up to 15 internal hard drives	Up to six DAE2 with a maximum of 15 hard drives each	Two per DPE
CX700 SPE	One DAE2-OS with at least five and up to 15 hard drives	Up to 15 DAE2 with a maximum of 15 hard drives each	Two per SPE and DAE2-OS



NOTE: The DAE2-OS is the first DAE2 enclosure that is connected to the CX700 and has the core software preinstalled on the first five hard drives.

Other Documents You May Need



The *Product Information Guide* provides important safety and regulatory information. Warranty information may be included within this document or as a separate document.



NOTE: To configure Dell blade server modules in a Dell PowerEdge Cluster, see the *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* document located on the Dell Support website at support.dell.com.

- The *Platform Guide* provides information about the platforms that support your cluster configuration.
- The *Rack Installation Guide* included with your rack solution describes how to install your system into a rack.
- The *Getting Started Guide* provides an overview of initially setting up your system.
- The *Installation and Troubleshooting Guide* describes how to troubleshoot the system and install or replace system components.
- The HBA documentation provides installation instructions for the HBAs.
- Systems management software documentation describes the features, requirements, installation, and basic operation of the software.
- Operating system documentation describes how to install (if necessary), configure, and use the operating system software.
- Documentation for any components you purchased separately provides information to configure and install those options.
- The Dell PowerVault™ tape library documentation provides information for installing, troubleshooting, and upgrading the tape library.
- The RAID documentation provides information for installing and configuring a RAID controller card.
- The documentation that came with your storage system.

- The EMC PowerPath documentation that came with your HBA kit(s).
- Updates are sometimes included with the system to describe changes to the system, software, and/or documentation.



NOTE: Always read the updates first because they often supersede information in other documents.

- Release notes or readme files may be included to provide last-minute updates to the system or documentation, or advanced technical reference material intended for experienced users or technicians.

Cabling Your Cluster Hardware



NOTE: To configure Dell blade server modules in a Dell PowerEdge Cluster, see the *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* document located on the Dell Support website at support.dell.com.

Cabling the Mouse, Keyboard, and Monitor

When installing a cluster configuration in a rack, you must include a switch box to connect the mouse, keyboard, and monitor to the nodes. See the documentation included with your rack for instructions on cabling each node's connections to the switch box.

Cabling the Power Supplies

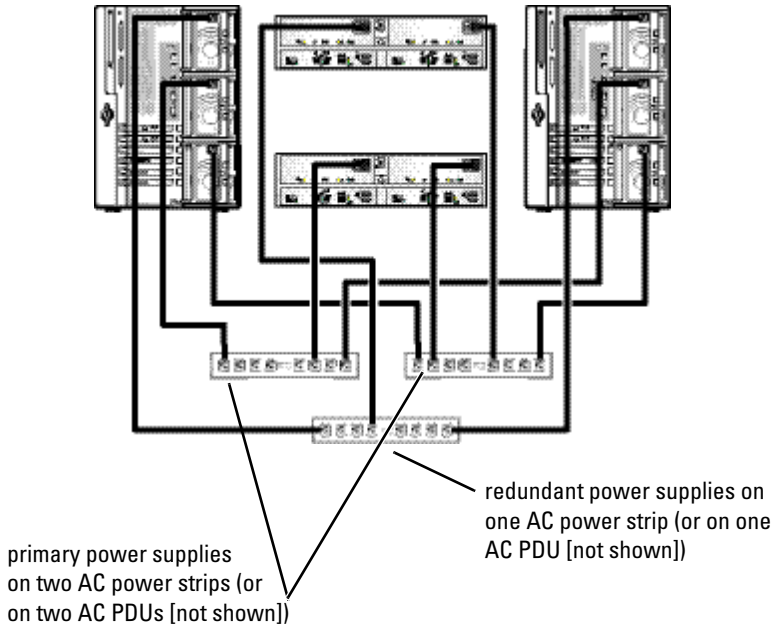
Refer to the documentation for each component in your cluster solution to ensure that the specific power requirements are satisfied.

The following guidelines are recommended to protect your cluster solution from power-related failures:

- For nodes with multiple power supplies, plug each power supply into a separate AC circuit.
- Use uninterruptible power supplies (UPS).
- For some environments, consider having backup generators and power from separate electrical substations.


Figure 2-1, Figure 2-2, and Figure 2-3 illustrate recommended methods for cabling power for a cluster solution consisting of two PowerEdge systems and two storage systems. To ensure redundancy, the primary power supplies of all the components are grouped onto one or two circuits and the redundant power supplies are grouped onto a different circuit.

Figure 2-1. Power Cabling Example With Three Power Supplies in the PowerEdge Systems



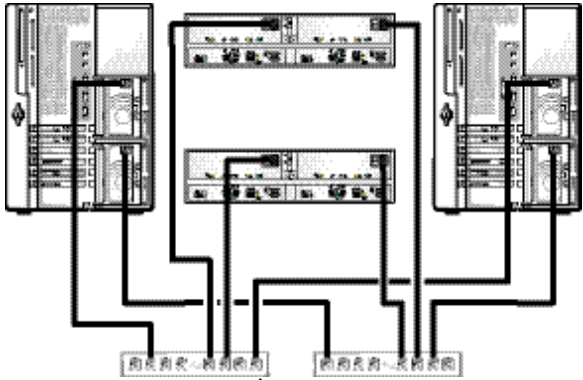
➔ NOTICE: This illustration is intended only to demonstrate the power distribution of the components.

Figure 2-2. Power Cabling Example With One Power Supply in the PowerEdge Systems



➡ NOTICE: This illustration is intended only to demonstrate the power distribution of the components.

Figure 2-3. Power Cabling Example With Two Power Supplies in the PowerEdge Systems



primary power supplies on one AC power strip (or on one AC PDU [not shown])

redundant power supplies on one AC power strip (or on one AC PDU [not shown])

NOTICE: This illustration is intended only to demonstrate the power distribution of the components.

Cabling Your Cluster For Public and Private Networks

NOTE: To configure Dell blade server modules in a Dell PowerEdge Cluster, see the *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* document located on the Dell Support website at support.dell.com.


The network adapters in the cluster nodes provide at least two network connections for each node, as described in Table 2-1.

Table 2-1. Network Connections

Network Connection	Description
Public network	All connections to the client LAN. At least one public network must be configured for Mixed mode for private network failover.
Private network	A dedicated connection for sharing cluster health and status information only.

Figure 2-4 shows an example of cabling in which dedicated network adapters in each node are connected to each other (for the private network) and the remaining network adapters are connected to the public network.

Figure 2-4. Example of Network Cabling Connection



Cabling the Public Network

Any network adapter supported by a system running TCP/IP may be used to connect to the public network segments. You can install additional network adapters to support additional public network segments or to provide redundancy in the event of a faulty primary network adapter or switch port.

Cabling the Private Network

The private network connection to the nodes is provided by a different network adapter in each node. This network is used for intra-cluster communications. Table 2-2 describes three possible private network configurations.

Table 2-2. Private Network Hardware Components and Connections

Method	Hardware Components	Connection
Network switch	Fast Ethernet or Gigabit Ethernet network adapters and switches	Connect <i>standard</i> Ethernet cables from the network adapters in the nodes to a Fast Ethernet or Gigabit Ethernet switch.
Point-to-Point Fast Ethernet (two-node clusters only)	Fast Ethernet network adapters	Connect a <i>crossover</i> Ethernet cable between the Fast Ethernet network adapters in both nodes.
Point-to-Point Gigabit Ethernet (two-node clusters only)	Copper Gigabit Ethernet network adapters	Connect a <i>standard</i> Ethernet cable between the Gigabit Ethernet network adapters in both nodes.



NOTE: On certain Microsoft® Windows® 2000 Advanced Server configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. See Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 at www.microsoft.com for more information. This issue has been corrected in Windows Server 2003.

Using Dual-Port Network Adapters

You can configure your cluster to use the public network as a failover for private network communications. If dual-port network adapters are used, do not use both ports simultaneously to support both the public and private networks.

NIC Teaming

NIC teaming combines two or more NICs to provide load balancing and fault tolerance. Your cluster supports NIC teaming, but only in a public network; NIC teaming is not supported in a private network.

You should use the same brand of NICs in a team, and you cannot mix brands of teaming drivers.


Fibre Channel Cable Connectors

Your cluster configuration requires fiber optic connectors to connect the various Fibre Channel storage components to the cluster configuration. Your solution includes at least one of the following types of connectors:

- Duplex LC multimode fiber optic connector
- Duplex SC multimode fiber optic connector

The duplex LC multimode fiber optic connector (see Figure 2-5) is used to connect a Fibre Channel switch to a Fibre Channel HBA port on a cluster node or to a storage system. This type of connection requires fiber optic cables with duplex LC multimode fiber optic connectors.

Figure 2-5. Duplex LC Multimode Fiber Optic Connector



The duplex SC multimode fiber optic connector (see Figure 2-6) may be used to connect a Fibre Channel switch to a Dell PowerVault™ 132T or 136T tape library. This type of connection requires fiber optic cables with both duplex LC and SC multimode fiber optic connectors.

Figure 2-6. Duplex SC Multimode Fiber Optic Connector




Table 2-3 provides the hardware component applications for the duplex multimode fiber optic connectors.

Table 2-3. Cable Connector Applications

Hardware Component	Duplex Multimode Fiber Optic Connector
PowerVault tape library	SC or LC
Cluster node HBA	LC
CX300 DPE	LC
CX500 DPE	LC
CX700 SPE	LC
Fibre Channel switch	LC

The LC and SC connectors consist of two individual fiber optic connectors. Each connector is indexed and must be inserted and aligned properly in the GBIC or SFP module connector.

➡ NOTICE: When using duplex multimode fiber optic connectors, keep the covers on the connectors until you are ready to insert the connectors into the system. The covers protect the cables and connectors and prevent foreign particles from entering and damaging the connector.

Cabling the Storage Systems


This section provides information for connecting your cluster to a storage system in a direct-attached configuration, or to one or more storage systems in a SAN-attached configuration.

Cabling Storage for Your Direct-Attached Cluster

A direct-attached cluster configuration consists of redundant Fibre Channel HBA ports cabled directly to a Dell | EMC storage system. Direct-attached configurations are self-contained and do not share any physical resources with other server or storage systems outside of the cluster.

Figure 2-7 shows an example of a direct-attached, single cluster configuration with redundant HBA ports installed in each cluster node.

Figure 2-7. Direct-Attached Cluster Configuration



Cabling One Cluster to a Dell | EMC Storage System

Each cluster node attaches to the storage system using two fiber optic cables with duplex LC multimode connectors that attach to the HBA ports in the cluster nodes and the SP ports in the Dell | EMC storage system. These connectors consist of two individual fiber optic connectors with indexed tabs that must be aligned properly into the HBA ports and SP ports.

NOTICE: Do not remove the connector covers until you are ready to insert the connectors into the HBA port, SP port, or tape library port.

NOTE: The connections listed in this section are representative of one proven method of ensuring redundancy in the connections between the cluster nodes and the storage system. Other methods that achieve the same type of redundant connectivity may be acceptable.

Cabling a Two-Node Cluster to a CX300 or CX500 Storage System

- 1 Connect cluster node 1 to the storage system.
 - a Install a cable from cluster node 1 HBA port 0 to SP-A port FE 0.
 - b Install a cable from cluster node 1 HBA port 1 to SP-B port FE 0.
- 2 Connect cluster node 2 to the storage system.
 - a Install a cable from cluster node 2 HBA port 0 to SP-A port FE 1.
 - b Install a cable from cluster node 2 HBA port 1 to SP-B port FE 1.

Figure 2-8 and Figure 2-9 illustrate methods of cabling a two-node direct-attached cluster to a CX300 and CX500 storage system, respectively.

NOTE: The cables are connected to the storage processor ports in sequential order for illustrative purposes. While the available ports in your storage system may vary, HBA port 0 and HBA port 1 must be connected to SP-A and SP-B, respectively.

Figure 2-8. Cabling the Cluster Nodes to a CX300 Storage System





Figure 2-9. Cabling the Cluster Nodes to a CX500 Storage System



Cabling a Two-Node Cluster to a CX700 Storage System

- 1** Connect cluster node 1 to the storage system.
 - a** Install a cable from cluster node 1 HBA port 0 to SP-A port 0.
 - b** Install a cable from cluster node 1 HBA port 1 to SP-B port 0.
- 2** Connect cluster node 2 to the storage system.
 - a** Install a cable from cluster node 2 HBA port 0 to SP-A port 1.
 - b** Install a cable from cluster node 2 HBA port 1 to SP-B port 1.

Figure 2-10 illustrates a method of cabling a two-node direct-attached cluster to a CX700 storage system.


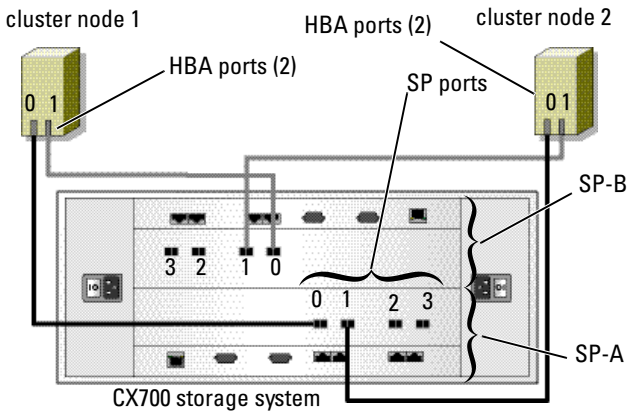
 **NOTE:** The cables are connected to the storage processor ports in sequential order for illustrative purposes. While the available ports in your storage system may vary, HBA port 0 and HBA port 1 must be connected to SP-A and SP-B, respectively.

Figure 2-10. Cabling the Cluster Nodes to a CX700 Storage System



Cabling a Four-Node Cluster to a CX300 and CX500 Storage System

The CX300 and CX500 storage systems do not support more than two cluster nodes in a direct-attached cluster configuration. Only the CX700 storage system can support a four-node direct-attached cluster configuration.

Cabling a Four-Node Cluster to a CX700 Storage System

- 1** Connect cluster node 1 to the storage system.
 - a** Install a cable from cluster node 1 HBA port 0 to SP-A port 0.
 - b** Install a cable from cluster node 1 HBA port 1 to SP-B port 0.
- 2** Connect cluster node 2 to the storage system.
 - a** Install a cable from cluster node 2 HBA port 0 to SP-A port 1.
 - b** Install a cable from cluster node 2 HBA port 1 to SP-B port 1.
- 3** Connect cluster node 3 to the storage system.
 - a** Install a cable from cluster node 3 HBA port 0 to SP-A port 2.
 - b** Install a cable from cluster node 3 HBA port 1 to SP-B port 2.
- 4** Connect cluster node 4 to the storage system.
 - a** Install a cable from cluster node 4 HBA port 0 to SP-A port 3.
 - b** Install a cable from cluster node 4 HBA port 1 to SP-B port 3.

Cabling Two Clusters to a Dell | EMC Storage System

The CX300 and CX500 storage systems do not support more than one direct-attached 2-node cluster.

The CX700 storage system includes four ports on each storage processor, allowing you to connect two 2-node clusters or a single four-node cluster running Windows Server 2003 to the storage system in a direct-attached configuration.



NOTE: EMC® Access Logix™ is required if the CX700 storage system is connected to more than one cluster in a direct-attached configuration.

Cabling Two 2-Node Clusters to a Dell | EMC CX700 Storage System

- 1** In the first cluster, connect cluster node 1 to the storage system.
 - a** Install a cable from cluster node 1 HBA port 0 to SP-A port 0.
 - b** Install a cable from cluster node 1 HBA port 1 to SP-B port 0.
- 2** In the first cluster, connect cluster node 2 to the storage system.
 - a** Install a cable from cluster node 2 HBA port 0 to SP-A port 1.
 - b** Install a cable from cluster node 2 HBA port 1 to SP-B port 1.
- 3** In the second cluster, connect cluster node 1 to the storage system.
 - a** Install a cable from cluster node 1 HBA port 0 to SP-A port 2.
 - b** Install a cable from cluster node 1 HBA port 1 to SP-B port 2.
- 4** In the second cluster, connect cluster node 2 to the storage system.
 - a** Install a cable from cluster node 2 HBA port 0 to SP-A port 3.
 - b** Install a cable from cluster node 2 HBA port 1 to SP-B port 3.

Cabling Storage for Your SAN-Attached Cluster

A SAN-attached cluster is a cluster configuration where all cluster nodes are attached to a single storage system or to multiple storage systems through a SAN using a redundant switch fabric.

SAN-attached cluster configurations provide more flexibility, expandability, and performance than direct-attached configurations.

See "Fibre Channel Switch Fabric" on page 16 for more information on Fibre Channel switch fabrics.

Figure 2-11 shows an example of a two node, SAN-attached cluster running Microsoft® Windows® 2000 Advanced Server.

Figure 2-12 shows an example of an eight-node, SAN-attached cluster running Windows Server 2003.

NOTE: The connections listed in this section are representative of one proven method of ensuring redundancy in the connections between the cluster nodes and the storage system. Other methods that achieve the same type of redundant connectivity may be acceptable.

Figure 2-11. SAN-Attached Cluster Running Windows 2000 Advanced Server


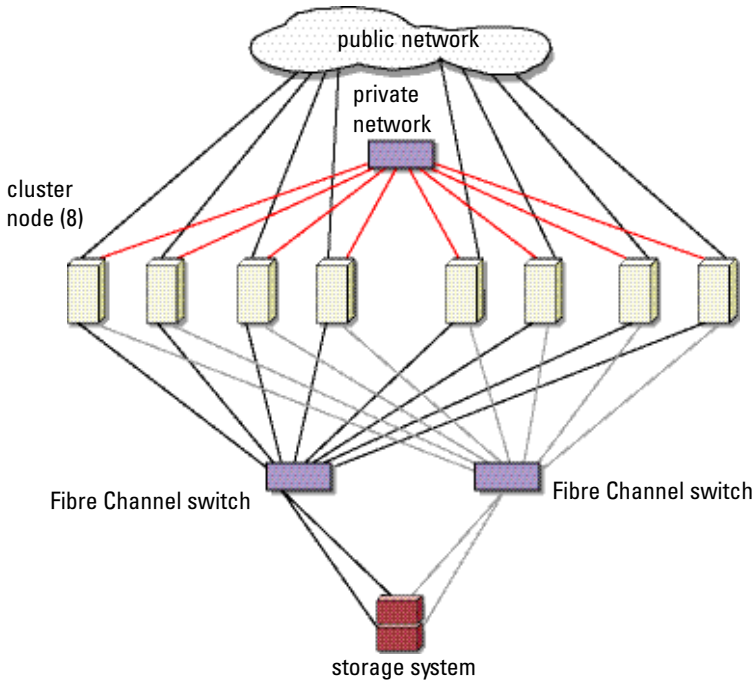


Figure 2-12. SAN-Attached Cluster Running Windows Server 2003



Cabling One SAN-Attached Cluster to a Dell | EMC Storage System

The supported Dell | EMC storage systems are configured as follows:

- CX300 — One DPE, one or more DAE2 enclosures (optional), and two SPSs
- CX500 — One DPE, one or more DAE2 enclosures (optional), and two SPSs
- CX700 — One SPE, at least one DAE2 (or DAE2-OS) enclosure, and two SPSs

The cluster nodes attach to the storage system using a redundant switch fabric and fiber optic cables with duplex LC multimode connectors.

The switches, the HBA ports in the cluster nodes, and the SP ports in the storage system use duplex LC multimode connectors. The connectors consist of two individual fiber optic connectors with indexed tabs that must be inserted and aligned properly in the small form-factor pluggable (SFP) module connectors on the Fibre Channel switches and the connectors on the cluster nodes and storage systems.

See "Cabling Your Cluster For Public and Private Networks" on page 26 for more information on the duplex LC multimode fiber optic connector.

Each HBA port is cabled to a port on a Fibre Channel switch. One to four cables connect from the outgoing ports on a switch to a storage processor on a Dell | EMC storage system.

Table 2-4 provides information for cabling your storage system to the Fibre Channel switch.

Figure 2-13 and Figure 2-14 illustrate methods for cabling a SAN-attached cluster to the CX300 and CX500 storage systems, respectively.

Figure 2-15 illustrates a method for cabling a SAN-attached cluster to a CX700 storage system.


 **NOTE:** The cables are connected to the storage processor ports in sequential order for illustrative purposes. While the available ports in your storage system may vary, HBA port 0 and HBA port 1 must be connected to SP-A and SP-B, respectively.


Table 2-4. Storage System Cabling Description

Storage System	SP Ports	Fiber Optic Cables Required	Cabling Description
CX300	Two ports per storage processor	4	Attach one cable from each storage processor port to the Fibre Channel switch.
CX500	Two ports per storage processor	4	
CX700	Four ports per storage processor	8	

Cabling a SAN-Attached Cluster to a Dell | EMC CX300 Storage System

- 1 Connect cluster node 1 to the SAN.
 - a Connect a cable from HBA port 0 to Fibre Channel switch 0 (sw0).
 - b Connect a cable from HBA port 1 to Fibre Channel switch 1 (sw1).
- 2 Repeat step 1 for each cluster node.
- 3 Connect the storage system to the SAN.
 - a Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port FE 1.
 - b Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port FE 0.
 - c Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port FE 0.
 - d Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port FE 1.

Figure 2-13. Cabling a SAN-Attached Cluster to the Dell | EMC CX300 DPE

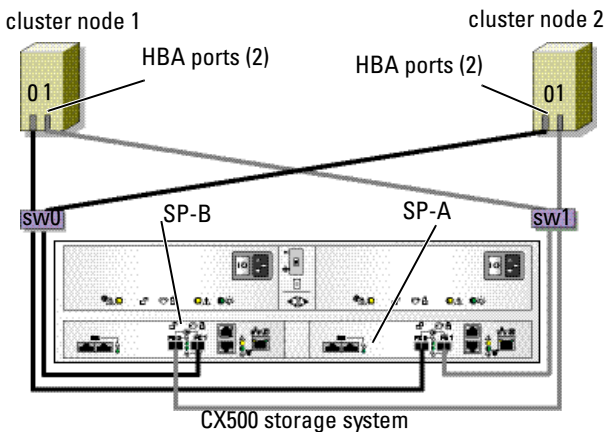


Cabling a SAN-Attached Cluster to a Dell | EMC CX500 Storage System

- 1 Connect cluster node 1 to the SAN.
 - a Connect a cable from HBA port 0 to Fibre Channel switch 0 (sw0).
 - b Connect a cable from HBA port 1 to Fibre Channel switch 1 (sw1).
- 2 Repeat step 1 for each node.

- 3** Connect the storage system to the SAN.
 - a** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port FE 0.
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port FE 1.
 - c** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port FE 1.
 - d** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port FE 0.

Figure 2-14. Cabling a SAN-Attached Cluster to the CX500 DPE

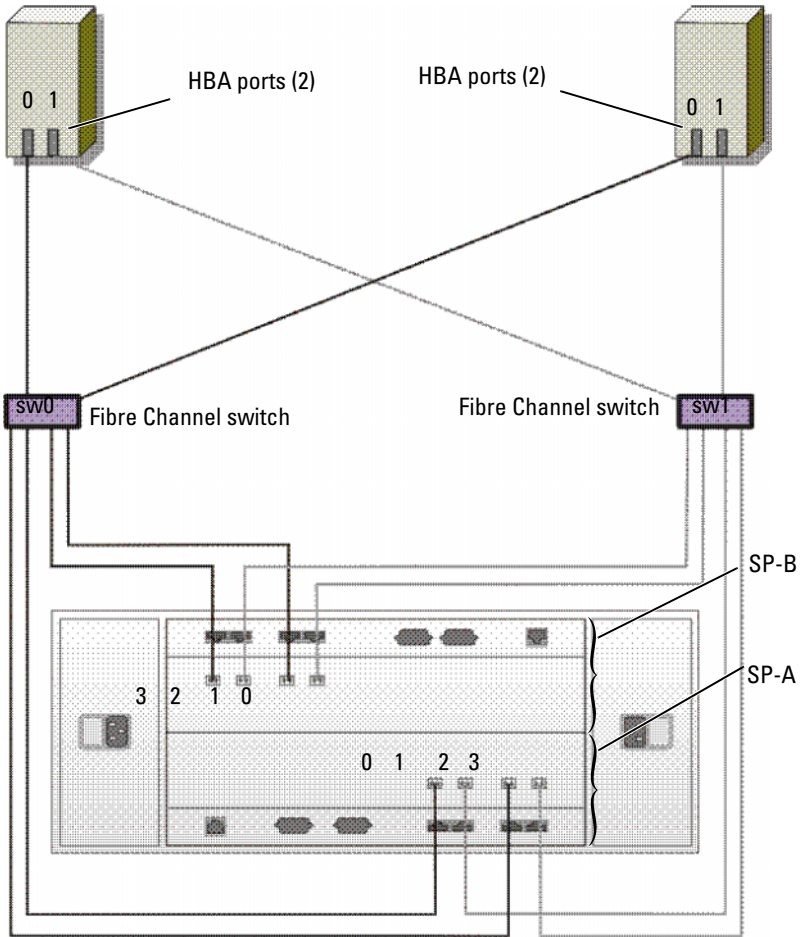


Cabling a SAN-Attached Cluster to the CX700 Storage System

- 1** Connect cluster node 1 to the SAN.
 - a** Connect a cable from HBA port 0 to Fibre Channel switch 0 (sw0).
 - b** Connect a cable from HBA port 1 to Fibre Channel switch 1 (sw1).
- 2** Repeat step 1 for each node.
- 3** Connect the storage system to the SAN.
 - a** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port 0.
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port 2.
 - c** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port 1.
 - d** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port 3.
 - e** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port 1.

- f Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port 3.
- g Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port 0.
- h Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port 2.

Figure 2-15. Cabling a SAN-Attached Cluster to the CX700 SPE



Cablings Multiple SAN-Attached Clusters to a Dell | EMC Storage System

To cable multiple clusters to the storage system, connect the cluster nodes to the appropriate Fibre Channel switches and then connect the Fibre Channel switches to the appropriate storage processors on the processor enclosure.

See the *Platform Guide* for rules and guidelines for SAN-attached clusters.



NOTE: The following procedures use Figure 2-13, Figure 2-14, and Figure 2-15 as examples for cabling additional clusters.

Cablings Multiple SAN-Attached Clusters to the CX300 Storage System

- 1 In the first cluster, connect cluster node 1 to the SAN.
 - a Connect a cable from HBA0 to Fibre Channel switch 0 (sw0).
 - b Connect a cable from HBA1 to Fibre Channel switch 1 (sw1).
- 2 In the first cluster, repeat step 1 for each node.
- 3 For each additional cluster, repeat step 1 and step 2.
- 4 Connect the storage system to the SAN.
 - a Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port FE 0.
 - b Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port FE 1.
 - c Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port FE 1.
 - d Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port FE 0.

Cablings Multiple SAN-Attached Clusters to the CX500 Storage System

- 1 In the first cluster, connect cluster node 1 to the SAN.
 - a Connect a cable from HBA port 0 to Fibre Channel switch 0 (sw0).
 - b Connect a cable from HBA port 1 to Fibre Channel switch 1 (sw1).
- 2 In the first cluster, repeat step 1 for each node.
- 3 For each additional cluster, repeat step 1 and step 2.
- 4 Connect the storage system to the SAN.
 - a Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port FE 0.
 - b Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port FE 1.
 - c Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port FE 1.
 - d Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port FE 0.

Cabling Multiple SAN-Attached Clusters to the CX700 Storage System

- 1** In the first cluster, connect cluster node 1 to the SAN.
 - a** Connect a cable from HBA port 0 to Fibre Channel switch 0 (sw0).
 - b** Connect a cable from HBA port 1 to Fibre Channel switch 1 (sw1).
- 2** In the first cluster, repeat step 1 for each node.
- 3** For each additional cluster, repeat step 1 and step 2.
- 4** Connect the storage system to the SAN.
 - a** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port 0.
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A port 2.
 - c** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port 1.
 - d** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B port 3.
 - e** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port 1.
 - f** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A port 3.
 - g** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port 0.
 - h** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B port 2.

Zoning Your Dell | EMC Storage System in a Switched Environment

Dell only supports single-initiator zoning for connecting clusters to a Dell | EMC storage system in a switched environment. When using EMC PowerPath, a separate zone is created from each HBA port to the DPE or SPE.

Connecting a PowerEdge Cluster to Multiple Storage Systems

You can increase your cluster storage capacity by attaching multiple storage systems to your cluster using a redundant switch fabric. PowerEdge cluster systems can support configurations with multiple storage units attached to clustered servers. In this scenario, the MSCS software can fail over disk drives in any cluster-attached shared storage array between the cluster nodes.


When attaching multiple storage systems with your cluster, the following rules apply:

- There is a maximum of four storage systems per cluster.
- The shared storage systems and firmware must be identical. Using dissimilar storage systems and firmware for your shared storage is not supported.
- MSCS is limited to 22 drive letters. Because drive letters A through D are reserved for local disks, a maximum of 22 drive letters (E to Z) can be used for your storage system disks.
- Windows Server 2003 supports mount points, allowing greater than 22 drives per cluster.

See "Assigning Drive Letters and Mount Points" on page 69 for more information.

Figure 2-16 provides an example of cabling the cluster nodes to four Dell | EMC storage systems. See "Fibre Channel Switch Fabric" on page 16 for more information.

Figure 2-16. PowerEdge Cluster Nodes Cabled to Four Storage Systems



Connecting a PowerEdge Cluster to a Tape Library

To provide additional backup for your cluster, you can add tape backup devices to your cluster configuration. The Dell PowerVault™ tape libraries contain an integrated Fibre Channel bridge, or Storage Network Controller (SNC), that connects directly to your Dell | EMC Fibre Channel switch.

Figure 2-17 shows a supported PowerEdge cluster configuration using redundant Fibre Channel switches and a tape library. In this configuration, each of the cluster nodes can access the tape library to provide backup for your local disk resources, as well as your cluster disk resources. Using this configuration allows you to add more servers and storage systems in the future, if needed.



 **NOTE:** While tape libraries can be connected to multiple fabrics, they do not provide path failover.

Figure 2-17. Cabling a Storage System and a Tape Library



Obtaining More Information


See the storage and tape backup documentation for more information on configuring these components.

Configuring Your Cluster With SAN Backup

You can provide centralized backup for your clusters by sharing your SAN with multiple clusters, storage systems, and a tape library.

Figure 2-18 provides an example of cabling the cluster nodes to your storage systems and SAN backup with a tape library.

Figure 2-18. Cluster Configuration Using SAN-Based Backup



Preparing Your Systems for Clustering

Before You Begin

- 1 Ensure that your site can handle the cluster's power requirements.
Contact your sales representative for information about your region's power requirements.



CAUTION: Only trained service technicians are authorized to remove and access any of the components inside the system. See your *Product Information Guide* for complete information about safety precautions, working inside the computer, and protecting against electrostatic discharge.

- 2 Ensure that the following components are installed in each PowerEdge system:
 - Network adapters
 - HBA ports
 - Hard drives
 - Any additional peripheral components
- 3 Configure the storage system(s) as described in your storage system documentation.
- 4 Cable the system hardware.
See "Cabling Your Cluster Hardware" on page 23 for more information.

Installation Overview

This section provides installation overview procedures for configuring a cluster running the Microsoft® Windows® 2000, Advanced Server or Windows Server™ 2003 operating system.

- 1 Ensure that the cluster meets the requirements as described in "Before You Begin" on page 47.

- 2 Select a domain model that is appropriate for the corporate network and operating system.

See "Selecting a Domain Model" on page 49.

- 3 Reserve static IP addresses for the cluster resources and components, including:

- Public network
- Private network
- Cluster virtual servers

Use these IP addresses when you install the Windows operating system and MSCS.

See "Installing the Fibre Channel HBAs" on page 57.

- 4 Configure the internal hard drives.

See "Configuring Internal Drives in the Cluster Nodes" on page 50.

- 5 Install and configure the Windows operating system.

The Windows operating system must be installed on all of the nodes. Each node must have a licensed copy of the Windows operating system, and a Certificate of Authenticity.

See "Installing and Configuring the Microsoft Windows Operating System" on page 50.

- 6 Install or update the HBA drivers.

See "Installing the Fibre Channel HBA Drivers" on page 57. For the specific HBA drivers required for the operating system, see the *Platform Guide*.

- 7 Install and configure the storage management software.

See the documentation included with the Dell | EMC storage system or available at the Dell Support website at support.dell.com.

- 8 Configure the hard drives on the shared storage system(s).

See "Configuring and Managing LUNs" on page 68.

- 9 Configure the MSCS software.

See "Installing and Configuring MSCS" on page 74.

- 10** Verify cluster functionality. Ensure that:
- The cluster components are communicating properly.
 - MSCS is started.

See "Verifying Cluster Functionality" on page 80.

- 11** Verify cluster resource availability.

Use Cluster Administrator to check the running state of each resource group. See "Verifying Cluster Resource Availability" on page 80.

The following subsections provide detailed information for each step in the "Installation Overview" on page 47 that is specific to the operating system.

Selecting a Domain Model

On a cluster running the Microsoft® Windows® operating system, all nodes must belong to a common domain or directory model. The following configurations are supported:

- All nodes are member servers in an Active Directory® domain.
- All nodes are domain controllers in an Active Directory domain.
- At least one node is a domain controller in an Active Directory and the remaining nodes are member servers.

Configuring the Nodes as Domain Controllers

If a node is configured as a domain controller, client system access to its cluster resources can continue even if the node cannot contact other domain controllers. However, domain controller functions can cause additional overhead, such as log on, authentication, and replication traffic.

If a node is not configured as a domain controller and the node cannot contact a domain controller, the node cannot authenticate client system requests.

Configuring Internal Drives in the Cluster Nodes

If your system uses a hardware-based RAID solution and you have added new internal hard drives to your system, or you are setting up the RAID configuration for the first time, you must configure the RAID array using the RAID controller's BIOS configuration utility before installing the operating system.

For the best balance of fault tolerance and performance, use RAID 1. See the RAID controller documentation for more information on RAID configurations.



NOTE: If you are not using a hardware-based RAID solution, use the Microsoft® Windows® Disk Management tool to provide software-based redundancy.

Installing and Configuring the Microsoft Windows Operating System



NOTICE: Windows standby mode and hibernation mode are not supported in cluster configurations. Do not enable either mode.

- 1 Ensure that the cluster configuration meets the requirements listed in "Before You Begin" on page 47.
- 2 Cable the hardware.



NOTE: Do not connect the nodes to the shared storage systems yet. See "Cabling Your Cluster Hardware" on page 23.

- 3 Install and configure the Windows 2000 Advanced Server operating system with the latest service pack or the Windows Server 2003 operating system with the latest service pack on each node.
See the *Platform Guide* for more information about the latest supported service pack.
- 4 If you are installing Windows Server 2003, go to step 5.
If you are installing Windows 2000 Advanced Server, select the option to install the Cluster Service files when you are prompted. You will configure the Cluster Service later.
- 5 Ensure that the network adapter drivers installed on each cluster node are the latest supported version.

- 6** Configure the public and private network adapter interconnects in each node, and place the interconnects on separate IP subnetworks using static IP addresses. See "Configuring Windows Networking" on page 52.
See the *Platform Guide* for information on required drivers.
- 7** Shut down both nodes and connect each node to shared storage.
See "Cabling Your Cluster Hardware" on page 23.
- 8** If required, configure the storage software.
- 9** Reboot node 1.
- 10** From node 1, write the disk signature and then partition, format, and assign drive letters and volume labels to the hard drives in the storage system using the Windows Disk Management application.
See "Naming and Formatting Drives on the Shared Storage System" on page 69.
- 11** On node 1, verify disk access and functionality on all shared disks.
- 12** Shut down node 1.
- 13** Verify disk access by performing the following steps on the other node:
 - a** Turn on the node.
 - b** Modify the drive letters to match the drive letters on node 1.
This procedure allows the Windows operating system to mount the volumes.
 - c** Close and reopen Disk Management.
 - d** Verify that Windows can see the file systems and the volume labels.
- 14** Turn on node 1.
- 15** Install and configure the Cluster Service.
See "Installing and Configuring Microsoft Cluster Service (MSCS) with Windows 2000" on page 74 and "Configuring Microsoft Cluster Service (MSCS) with Windows Server 2003" on page 75.
- 16** Install and set up the application programs (optional).
- 17** Enter the cluster configuration information on the "Cluster Data Form" on page 129 (optional).

Configuring Windows Networking

You must configure the public and private networks in each node before you install MSCS. The following subsections introduce you to some principles and procedures necessary for the networking prerequisites.

Assigning Static IP Addresses to Cluster Resources and Components

A static IP address is an Internet address that a network administrator assigns exclusively to a system or a resource. The address assignment remains in effect until it is changed by the network administrator.

The IP address assignments for the cluster's public LAN segments depend on the environment's configuration. Configurations running the Windows operating system require static IP addresses assigned to hardware and software applications in the cluster, as listed in Table 3-1.

Table 3-1. Applications and Hardware Requiring IP Address Assignments

Application/Hardware	Description
Cluster IP address	The cluster IP address is used for cluster management and must correspond to the cluster name. Because each server has at least two network adapters, the minimum number of static IP addresses required for a cluster configuration is five (one for each network adapter and one for the cluster). Additional static IP addresses are required when MSCS is configured with application programs that require IP addresses, such as file sharing.
Cluster-aware applications running on the cluster	These applications include Microsoft SQL Server, Enterprise Edition Microsoft Exchange Server, and Internet Information Server (IIS). For example, Microsoft SQL Server, Enterprise Edition requires at least one static IP address for the virtual server (Microsoft SQL Server does not use the cluster's IP address). Also, each IIS Virtual Root or IIS Server instance configured for failover needs a unique static IP address.

Table 3-1. Applications and Hardware Requiring IP Address Assignments (continued)

Application/Hardware	Description
Cluster node network adapters	<p>For cluster operation, two network adapters are required: one for the public network (LAN/WAN) and another for the private network (sharing heartbeat information between the nodes).</p> <p>See "Cabling Your Cluster Hardware" on page 23 for more information about cluster interconnect options.</p> <p>NOTE: To ensure operation during a DHCP server failure, use static IP addresses.</p> <p>NOTE: On certain Windows 2000 Advanced Server configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. See Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 at www.microsoft.com for more information. This issue has been corrected in Windows Server 2003.</p>

Configuring IP Addresses for the Private Network

Use the static IP address assignments for the network adapters used for the private network (cluster interconnect).




NOTE: The IP addresses in Table 3-2 are used as examples only.

Table 3-2. Examples of IP Address Assignments

Usage	Cluster Node 1	Cluster Node 2
Public network static IP address (for client and domain controller communications)	192.168.1.101	192.168.1.102
Public network subnet mask	255.255.255.0	255.255.255.0
Default gateway	192.168.1.1	192.168.1.1
WINS servers	Primary 192.168.1.11	Primary 192.168.1.11
	Secondary 192.168.1.12	Secondary 192.168.1.12


Table 3-2. Examples of IP Address Assignments (continued)

Usage	Cluster Node 1	Cluster Node 2
DNS servers	Primary 192.168.1.21	Primary 192.168.1.21
	Secondary 192.168.1.22	Secondary 192.168.1.22
Private network static IP address cluster interconnect (for node-to-node communications)	10.0.0.1	10.0.0.2
Private network subnet mask	255.255.255.0	255.255.255.0

 **NOTE:** Do not configure Default Gateway, NetBIOS, WINS, and DNS on the private network. If you are running Windows 2000 Advanced Server or Windows Server 2003, disable NetBIOS on the private network.

If multiple cluster interconnect network adapters are connected to a network switch, ensure that all of the private network's network adapters have a unique address. You can continue the IP address scheme in Table 3-2 with 10.0.0.3, 10.0.0.4, and so on for the private network's network adapters or network adapter teams of the other clusters connected to the same switch.

You can improve fault tolerance by using network adapters that support adapter teaming or by having multiple LAN segments. To avoid communication problems, do not use dual-port network adapters for the cluster interconnect.

 **NOTE:** NIC teaming is supported only on a public network, not on a private network.

Creating Separate Subnets for the Public and Private Networks

The public and private network's network adapters installed in the same cluster node must reside on separate IP subnetworks. Therefore, the private network used to exchange heartbeat information between the nodes must have a separate IP subnet or a different network ID than the public network, which is used for client connections.

Setting the Network Interface Binding Order for Clusters Running Windows 2000

- 1** On the Windows 2000 desktop, right-click **My Network Places**, and then click **Properties**.

The **Network and Dial-up Connections** window appears.

- 2** Click the **Advanced** menu and then click **Advanced Settings**.

The **Advanced Settings** window appears.

- 3** In the **Adapters and Bindings** tab in the **Connections** box, ensure that the **Public** connections designated for **Client access only** or **All communications** are at the top of the list.

To change the connection order:

- a** Click **Public** or **Private**.
- b** Click the up-arrow or down-arrow to move the connection type to the top or bottom of the **Connections** box.
- c** Click **OK**.
- d** Close the **Network and Dial-up Connections** window.

Setting the Network Interface Binding Order for Clusters Running Windows Server 2003

- 1** Click the **Start** button, select **Control Panel**, and double-click **Network Connections**.

- 2** Click the **Advanced** menu, and then click **Advanced Settings**.

The **Advanced Settings** window appears.

- 3** In the **Adapters and Bindings** tab, ensure that the **Public** connection is at the top of the list and followed by the **Private** connection.

To change the connection order:

- a** Click **Public** or **Private**.
- b** Click the up-arrow or down-arrow to move the connection to the top or bottom of the **Connections** box.
- c** Click **OK**.
- d** Close the **Network Connections** window.

Dual-Port Network Adapters and Adapter Teams in the Private Network

Dual-port network adapters and network adapter teams are not supported in the private network. They are supported only in the public network.

Verifying Communications Between Nodes

1 Open a command prompt on each cluster node.

2 At the prompt, type:

```
ipconfig /all
```

3 Press <Enter>.

All known IP addresses for each local server appear on the screen.

4 Issue the **ping** command from each remote system.

Ensure that each local server responds to the **ping** command. If the IP assignments are not set up correctly, the nodes may not be able to communicate with the domain. See "Troubleshooting" for more information.

Configuring the Internet Connection Firewall

Microsoft Windows Server 2003, Enterprise x64 Edition and Windows Server 2003, Enterprise Edition Service Pack 1 include an enhanced Internet Connection Firewall that can be configured to block incoming network traffic to a Poweredge system. To prevent the Internet Connection Firewall from disrupting cluster communications, additional configuration settings are required for Poweredge systems that are configured as cluster nodes in an MSCS cluster.

Certain network communications are necessary for cluster operations, for applications and services hosted by the cluster, and for clients accessing those services. If the Internet Connection Firewall is enabled on the cluster nodes, install and run the Security Configuration Wizard and then configure access for the cluster service and for any applications or services hosted by the cluster and the operating system.

See the following Microsoft Knowledge Base articles located at the Microsoft Support website at support.microsoft.com for more information:

- KB883398 - Internet Connection Firewall
- KB832017 - Network ports used by the Windows Server 2003 operating system

Installing the Fibre Channel HBAs

For dual HBA configurations, Dell recommends installing Fibre Channel HBAs on separate PCI buses. Placing the adapters on separate buses improves availability and performance.

See the *Platform Guide* for more information about your system's PCI bus configuration and supported HBAs.

Installing the Fibre Channel HBA Drivers

See the EMC documentation that is included with your HBA kit for more information.

See the Emulex support website located at www.emulex.com or the Dell Support website at support.dell.com for information about installing and configuring Emulex HBAs and EMC approved drivers.

See the QLogic support website at www.qlogic.com or the Dell Support website at support.dell.com for information about installing and configuring QLogic HBAs and EMC approved drivers.

See the *Platform Guide* for information about supported HBA controllers and drivers.

Implementing Zoning on a Fibre Channel Switched Fabric

A Fibre Channel switched fabric consists of one or more Fibre Channel switches that provide high-speed connections between servers and storage devices. The switches in a Fibre Channel fabric provide a connection through inbound and outbound points from one device (sender) to another device or switch (receiver) on the network. If the data is sent to another switch, the process repeats itself until a connection is established between the sender and the receiver.

Fibre Channel switches provide you with the ability to set up barriers between different devices and operating environments. These barriers create logical fabric subsets with minimal software and hardware intervention. Similar to subnets in the client/server network, logical fabric subsets divide a fabric into similar groups of components, regardless of their proximity to one another. The logical subsets that form these barriers are called *zones*.

Zoning automatically and transparently enforces access of information to the zone devices. More than one PowerEdge cluster configuration can share Dell | EMC storage system(s) in a switched fabric using Fibre Channel switch zoning and Access Logix. By using Fibre Channel switches to implement zoning, you can segment the SANs to isolate heterogeneous servers and storage systems from each other.

Using Zoning in SAN Configurations Containing Multiple Hosts

Using the combination of zoning and Access Logix in SAN configurations containing multiple hosts, you can restrict server access to specific volumes on a shared storage system by preventing the hosts from discovering a storage volume that belongs to another host. This configuration allows multiple clustered or nonclustered hosts to share a storage system.

Using Worldwide Port Name Zoning

PowerEdge cluster configurations support worldwide port name zoning.

A worldwide name (WWN) is a unique numeric identifier assigned to Fibre Channel interfaces, such as HBA ports, SP ports, and Fibre Channel to SCSI bridges or SNCs.

A WWN consists of an 8-byte hexadecimal number with each byte separated by a colon. For example, 10:00:00:60:69:00:00:8a is a valid WWN. Using WWN port name zoning allows you to move cables between switch ports within the fabric without having to update the zones.

Table 3-3 provides a list of WWN identifiers that you can find in the Dell | EMC cluster environment.

Table 3-3. Port Worldwide Names in a SAN Environment

Identifier	Description
xx : xx : 00 : 60 : 69 : xx : xx : xx	Dell EMC or Brocade switch
xx : xx : xx : 00 : 88 : xx : xx : xx	McData switch
50 : 06 : 01 : 6x : xx : xx : xx : xx	Dell EMC storage processor

Table 3-3. Port Worldwide Names in a SAN Environment (continued)

Identifier	Description
xx:xx:00:00:C9:xx:xx :xx	Emulex HBA ports
xx:xx:00:E0:8B:xx:xx :xx	QLogic HBA ports (non-embedded)
xx:xx:00:0F:1F:xx:xx :xx	QLA2362M HBA port
xx:xx:xx:60:45:xx:xx :xx	PowerVault 132T and 136T tape libraries
xx:xx:xx:E0:02:xx:xx :xx	PowerVault 128T tape autoloader
xx:xx:xx:C0:01:xx:xx :xx	PowerVault 160T tape library and Fibre Channel tape drives



NOTICE: When you replace a Fibre Channel HBA in a PowerEdge server, reconfigure your zones to provide continuous client data access. Additionally, when you replace a switch module, reconfigure your zones to prevent data loss or corruption.



NOTICE: You must configure your zones before you configure the LUNs and storage groups. Failure to do so may cause data loss, data corruption, or data unavailability.

Single Initiator Zoning

Each host HBA port in a SAN must be configured in a separate zone on the switch with the appropriate storage ports. This zoning configuration, known as *single initiator zoning*, prevents different hosts from communicating with each other, thereby ensuring that Fibre Channel communications between the HBAs and their target storage systems do not affect each other.

When you create your single-initiator zones, follow these guidelines:

- Create a zone for each HBA port and its target storage devices.
- Each CX300 storage processor can be connected to a maximum of 32 HBA ports in a SAN-attached environment.
- Each CX500 storage processor port can be connected to a maximum of 64 HBA ports in a SAN-attached environment.
- Each CX700 storage processor port can be connected to a maximum of 64 HBA ports in a SAN-attached environment.
- Each host can be connected to a maximum of four storage systems.
- The integrated bridge or SNC on a tape library can be added to any zone.



NOTE: If you are sharing a storage system with multiple clusters or a combination of clustered and nonclustered systems (hosts), you must enable EMC Access Logix™ and Access Control. Otherwise, you can only have one nonclustered system or one PowerEdge cluster attached to the Dell | EMC storage system.

Installing and Configuring the Shared Storage System

See "Shared Storage Systems" on page 11 for a list of supported Dell | EMC storage systems.

To install and configure the Dell | EMC storage system in your cluster:

- 1 Update the core software on your storage system and enable the EMC Access Logix software (optional) and install any additional software options, including EMC SnapView™, EMC MirrorView™, and SAN Copy™. See your EMC Navisphere documentation for more information.
- 2 Install the EMC Navisphere Agent™ and EMC PowerPath™ software on each cluster node.
See your Navisphere documentation for more information.
- 3 Update the storage system configuration settings using Navisphere Manager.
See "Enabling Access Logix and Creating Storage Groups Using Navisphere 6.x" on page 66 for more information.

The following subsections provide an overview of the storage management software and procedures for connecting the host systems to the storage systems.

Access Logix

Fibre Channel topologies allow multiple clusters and stand-alone systems to share a single storage system. However, if you cannot control access to the shared storage system, you can corrupt your data. To share your Dell | EMC storage system with multiple heterogeneous host systems and restrict access to the shared storage system, you can enable and configure the Access Logix software.

Access Logix is an optional software component that restricts LUN access to specific host systems. Using Access Logix software, you can:

- Connect multiple cluster nodes and stand-alone systems to a storage system.
- Create storage groups to simplify LUN management.
- Restrict LUN access to preassigned storage groups for data security.

Access Logix is enabled by configuring the Access Logix option on your storage system.

The storage systems are managed through a *management station*—a local or remote system that communicates with Navisphere Manager and connects to the storage system through an IP address. Using Navisphere Manager, you can secure your storage data by partitioning your storage system arrays into LUNs, assign the LUNs to one or more storage groups, and then restrict access to the LUNs by assigning the storage groups to the appropriate host systems.

Access Logix is required if:

- The server modules are configured in dissimilar configurations. These configurations include:
 - Two or more stand-alone systems/non-clustered hosts.
 - Two or more clusters.
 - Any combination of server modules configured as cluster nodes and stand-alone systems/non-clustered hosts.
- MirrorView, SnapView, or SAN Copy are installed on your attached storage system(s) and running in the cluster configuration.

Table 3-4 provides a list of cluster and host system configurations and their Access Logix requirement.

Table 3-4. Access Logix Software Requirements

Cluster Configuration	Access Logix Required
Single host	No
or	
One cluster	
Two or more clusters	Yes
or	
Two or more stand alone systems/non-clustered hosts	
or	
Any combination of clusters and non-clustered hosts	

Access Control

Access Control is a feature of Access Logix that connects the host system to the storage system. Enabling **Access Control** prevents all host systems from accessing any data on the storage system until they are given explicit access to a LUN through a storage group. By installing Access Logix on your storage system(s) and enabling **Access Control**, you can prevent the host systems from taking ownership of all LUNs on the storage system and prevent unauthorized access to sensitive information.

Access Control is enabled using Navisphere Manager. After you enable Access Logix and connect to the storage system from a management station, **Access Control** appears in the **Storage System Properties** window of Navisphere Manager. After you enable **Access Control** in Navisphere Manager, you are using Access Logix.

See "Dell | EMC Storage Management Software" on page 13 for additional information on Access Logix and Navisphere Manager.

After you enable **Access Control**, the host system can only read from and write to specific LUNs on the storage system. This organized group of LUNs and hosts is called a *storage group*.

Storage Groups

A storage group is a collection of one or more LUNs that are assigned to one or more host systems. Managed by Navisphere Manager, storage groups provide an organized method of assigning multiple LUNs to a host system. After you create LUNs on your storage system, you can assign the LUNs to a storage group in Navisphere Manager and then assign the storage group to a specific host. Because the host can only access its assigned storage group, it cannot access any LUNs assigned to other host systems, thereby protecting your data from unauthorized access.

To create the storage groups for your host systems, you must use Navisphere Manager and enable **Access Control** in the storage system.



NOTE: A host system can access only one storage group per storage system.

Table 3-5 describes the properties in the storage group.

Table 3-5. Storage Group Properties

Property	Description
Unique ID	A unique identifier that is automatically assigned to the storage group that cannot be changed.
Storage group name	The name of the storage group. The default storage group name is formatted as <i>Storage Group n</i> , where <i>n</i> equals the existing number of storage groups plus one.
Connected hosts	Lists the host systems connected to the storage group. Each host entry contains the following fields: <ul style="list-style-type: none">• Name — Name of the host system• IP address — IP address of the host system• OS — Operating system that is running on the host system NOTE: In a clustered environment, all nodes of a cluster must be connected to the same storage group.

Table 3-5. Storage Group Properties (continued)

Property	Description
Used host connection paths	<p>An additional storage group feature that performs the following tasks:</p> <ul style="list-style-type: none">• Lists all of the paths from the host server to the storage group• Displays whether the path is enabled or disabled <p>Each path contains the following fields:</p> <ul style="list-style-type: none">– HBA — Device name of the HBA in the host system– HBA Port — Unique ID for the HBA port connected to the storage system– SP Port — Unique ID for the storage processor port connected to the HBA port– SP ID — ID of the storage processor
LUNs in storage group	<p>Lists the LUNs in the storage group.</p> <p>Each LUN entry contains the following fields:</p> <ul style="list-style-type: none">• Identifier — LUN icon representing the LUN• Name — Name of the LUN• Capacity — Amount of allocated storage space on the LUN

Navisphere Manager

Navisphere Manager provides centralized storage management and configuration from a single management console. Using a GUI, Navisphere Manager allows you to configure and manage the disks and components in one or more shared storage systems.

You can access Navisphere Manager through a Web browser. Using Navisphere Manager, you can manage a Dell | EMC storage system either locally on the same LAN or through an Internet connection. Navisphere components (Navisphere Manager UI and Storage Management Server) are installed on a Dell | EMC storage system. You can access Navisphere Manager by opening a browser and entering the IP address of the storage system's SP. Navisphere Manager downloads components to your system and runs in the web browser.

Optionally, you can run Navisphere Management Server for Windows. This software component installs on a host system connected to a Dell | EMC storage system, allowing you to run Navisphere Storage Management Server on the host system.

Using Navisphere Manager, you can:

- Create storage groups for your host systems
- Create, bind, and unbind LUNs
- Change configuration settings
- Monitor storage systems

See "Dell | EMC Storage Management Software" on page 13 for more information on Navisphere Manager.

Navisphere Agent

Navisphere Agent is installed on the host system and performs the following tasks:


- Registers each host with the storage system
- Communicates configuration information from the host to the storage system


EMC PowerPath

PowerPath automatically reroutes Fibre Channel I/O traffic from the host system and a Dell | EMC CX-series storage system to any available path if a primary path fails for any reason. Additionally, PowerPath provides multiple path load balancing, allowing you to balance the I/O traffic across multiple SP ports.

Enabling Access Logix and Creating Storage Groups Using Navisphere 6.x

The following subsection provides the required procedures for creating storage groups and connecting your storage systems to the host systems using the Access Logix software.

 **NOTICE:** Before enabling **Access Control**, ensure that no hosts are attempting to access the storage system. Enabling **Access Control** prevents all hosts from accessing any data until they are given explicit access to a LUN in the appropriate storage group. You must stop all I/O before enabling **Access Control**. Dell recommends shutting down all hosts connected to the storage system during this procedure or data loss may occur. After you enable the **Access Control software**, it cannot be disabled.

- 1 Ensure that Navisphere Agent is started on all host systems.
 - a Click the **Start** button and select **Programs**→**Administrative Tools**, and then select **Services**.
 - b In the **Services** window, verify the following:
 - In the **Name** column, **Navisphere Agent** appears.
 - In the **Status** column, **Navisphere Agent** is set to **Started**.
 - In the **Startup Type** column, **Navisphere Agent** is set to **Automatic**.
- 2 Open a Web browser.
- 3 Enter the IP address of the storage management server on your storage system and then press <Enter>.  **NOTE:** The storage management server is usually one of the SPs on your storage system.
- 4 In the **Enterprise Storage** window, click the **Storage** tab.
- 5 Right-click the icon of your storage system.
- 6 In the drop-down menu, click **Properties**.
The **Storage Systems Properties** window appears.
- 7 Click the **Storage Access** tab.
- 8 Select the **Access Control Enabled** check box.
A dialog box appears, prompting you to enable **Access Control**.

- 9 Click **Yes** to enable **Access Control**.
- 10 Click **OK**.
- 11 Right-click the icon of your storage system and select **Create Storage Group**.

The **Create Storage Group** dialog box appears.

- 12 In the **Storage Group Name** field, enter a name for the storage group.

- 13 Click **Apply**.

- 14 Add new LUNs to the storage group.

- a Right-click the icon of your storage group and select **Properties**.
- b Click the **LUNs** tab.
- c In the **Available LUNs** window, click an available LUN.
- d Click the right-arrow button to move the selected LUN to the **Selected LUNs** pane.
- e Click **Apply**.

- 15 Add new hosts to the **Sharable** storage group.

- a In the **Storage Group Properties** dialog box, click the **Hosts** tab.
- b In the **Available Hosts** window pane, click the host system that you want to add to the storage group.
- c Using the right-arrow button, move the selected host to the **Hosts to be Connected** window pane.
- d Repeat step b and step c to add additional hosts.
- e Click **Apply**.

- 16 Click **OK** to exit the **Storage Group Properties** dialog box.

Configuring the Hard Drives on the Shared Storage System(s)

This section provides information for configuring the hard drives on the shared storage systems. The shared storage system hard drives must be configured before use. The following sections provide information on these configurations.

Configuring and Managing LUNs

Configuring and managing LUNs is accomplished using the Navisphere Manager utility. Before using Navisphere Manager, ensure that the Navisphere Agent service is started on your cluster nodes.

In some cases, the LUNs may have been bound when the system was shipped. It is still important, however, to install the management software and to verify that the desired LUN configuration exists.

You can manage your LUNs remotely using Navisphere Manager. A minimum of one LUN (RAID drive) is required for an active/passive configuration; at least two drives are required for an active/active configuration.

Dell recommends creating at least one LUN or virtual disk for each application. If multiple NTFS partitions are created on a single LUN or virtual disk, these partitions will not be able to fail over individually from node-to-node.

Using the Windows Dynamic Disks and Volumes

The Windows 2000 Advanced Server and Windows Server 2003 operating systems that are shipped from Microsoft do not support dynamic disks (upgraded disks) or volumes as shared storage in a cluster environment. If the shared cluster storage is configured as a dynamic disk, the Cluster Configuration wizard is not able to discover the disks, preventing the cluster and network clients from accessing the disks.

Configuring the RAID Level for the Shared Storage Subsystem

The hard drives in your shared storage subsystem must be configured into LUNs or virtual disks using Navisphere Manager. All LUNs or virtual disks, especially if they are used for the quorum resource, should be bound and incorporate the appropriate RAID level to ensure high availability.

See "Installing the Quorum Resource" on page 78 for more information on the quorum resource.



NOTE: It is recommended that you use a RAID level other than RAID 0 (which is commonly called striping). RAID 0 configurations provide very high performance, but do not provide the level of availability required for the quorum resource. See the documentation for your storage system for more information about setting up RAID levels for the system.

Naming and Formatting Drives on the Shared Storage System

When the LUNs have completed the binding process, assign drive letters to the LUNs and then format the drives as NTFS drives. Format the drives and assign volume labels from the first cluster node. When completed, the remaining nodes will see the file systems and volume labels.



NOTICE: Accessing the hard drives from multiple cluster nodes may corrupt the file system.

Assigning LUNs to Hosts

If you have **Access Control** enabled in Navisphere Manager, you must create storage groups and assign LUNs to the proper host systems.

Assigning Drive Letters and Mount Points

A mount point is a drive attached to an empty folder on an NTFS volume. A mount point drive functions the same as a normal drive, but is assigned a label or name instead of a drive letter. Using mount points, a cluster can support more shared disks than the number of available drive letters.

The cluster installation procedure does not automatically add the mount point into the disks managed by the cluster. To add the mount point to the cluster, create a physical disk resource in the cluster resource group for each mount point. Ensure that the new physical disk resource is in the same cluster resource group and is dependent on the root disk.



NOTE: Mount points are only supported in MSCS on the Windows Server 2003 operating system. When mounting a drive to an NTFS volume, do not create mount points from the quorum resource or between the clustered disks and the local disks. Mount points must be in the same cluster resource group and must be dependent on the root disk.



NOTICE: If the disk letters are manually assigned from the remaining node(s), the shared disks are simultaneously accessible from both nodes. To ensure file system integrity and prevent possible data loss before you install the MSCS software, prevent any I/O activity to the shared drives by performing this procedure on one node at a time, and ensure that all other nodes are shut down.

The number of drive letters required by individual servers in a cluster may vary. It is recommended that the shared drives be named in reverse alphabetical order beginning with the letter *z*.

To assign drive letters, create mount points, and format the disks on the shared storage system:

- 1 With the remaining node(s) shut down, open **Disk Management** on node 1.
- 2 Allow Windows to enter a signature on all new physical or logical drives.



NOTE: Do not create dynamic disks on your hard drives.

- 3 Locate the icon for the first unnamed, unformatted drive on the shared storage system.
- 4 Right-click the icon and select **Create** from the submenu.

If the unformatted drives are not visible, verify the following:

- The HBA driver is installed.
- The storage system is properly cabled to the servers.
- The LUNs and hosts are assigned through a storage group (if **Access Control** is enabled).

- 5 In the dialog box, create a partition the size of the entire drive (the default) and then click **OK**.




NOTE: The MSCS software allows only one node to access a logical drive at a time. If a logical drive is partitioned into multiple disks, only one node is able to access all the partitions for that logical drive. If each node needs to access a separate disk, two or more logical drives must be present in the storage system.

- 6 Click **Yes** to confirm the partition.
- 7 With the mouse pointer on the same icon, right-click and select **Change Drive Letter and Path** from the submenu.
- 8 Assign a drive letter to an NTFS volume or create a mount point.

To assign a drive letter to an NTFS volume:

- a Click **Edit** and select the letter you want to assign to the drive (for example, Z).
- b Click **OK**.
- c Go to step 9.

To create a mount point:

- a** Click **Add**.
 - b** Click **Mount in the following empty NTFS folder**.
 - c** Type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.
 - d** Click **OK**.
 - e** Go to step 9.
- 9** Click **Yes** to confirm the changes.
- 10** Right-click the drive icon again and select **Format** from the submenu.
- 11** Under **Volume Label**, enter a descriptive name for the new volume; for example, `Disk_Z` or `Email_Data`.
- 12** In the dialog box, change the file system to **NTFS**, select **Quick Format**, and click **Start**.
-  **NOTE:** The NTFS file system is required for shared-disk resources under MSCS.
- 13** Click **OK** at the warning.
- 14** Click **OK** to acknowledge that the format is complete.
- 15** Click **Close** to close the dialog box.
- 16** Repeat step 3 through step 15 for each remaining drive.
- 17** Close **Disk Management**.
- 18** Shut down node 1.
- 19** Perform the following steps on the remaining node(s), one at a time:
- a** Turn on the node.
 - b** Open **Disk Management**.
 - c** Assign the drive letters to the drives.
This procedure allows Windows to mount the volumes.
 - d** Reassign the drive letter, if necessary.
To reassign the drive letter, repeat step 7 through step 9.
 - e** Power down the node.

Configuring Hard Drive Letters When Using Multiple Shared Storage Systems

Before installing MSCS, ensure that both nodes have the same view of the shared storage systems. Because each node has access to hard drives that are in a common storage array, each node must have identical drive letters assigned to each hard drive. Using volume mount points in Windows Server 2003, your cluster can access more than 22 volumes.

The maximum number of accessible volumes depends on the operating system that is running in the cluster:

- Windows 2000 — Up to 22 logical drive letters (E through Z).
- Windows Server 2003 — Up to 22 logical drive letters plus additional volumes using volume mount points.

See "Assigning Drive Letters and Mount Points" on page 69 for more information.



NOTE: Drive letters A through D are reserved for the local system.

To ensure that hard drive letter assignments are identical:

- 1 Ensure that your cables are attached to the shared storage devices in the proper sequence.

You can view all of the storage devices using Windows 2000 or Windows Server 2003 Disk Management. Windows 2000 Disk Management displays all of the accessible disks from the first HBA, followed by all of the accessible disks from the second HBA. The disks attached to a port with a lower port number on the Fibre Channel switch are displayed first, followed by those with a higher port number.

- 2 To maintain proper drive letter assignments, ensure the first HBA detected by each node is connected to the first switch or SP-A and the second detected HBA is connected to the second switch or SP-B.

See "Cabling the Power Supplies" on page 23 for the location of SP-A and SP-B on the CX-series storage systems.

- 3 Go to "Formatting and Assigning Drive Letters and Volume Labels to the Disks" on page 73.

Formatting and Assigning Drive Letters and Volume Labels to the Disks

- 1** Shut down all the cluster nodes except node 1.
- 2** Format the disks, assign the drive letters and volume labels on node 1 by using the Windows Disk Management utility.
For example, create volumes labeled "Volume Y" for disk Y and "Volume Z" for disk Z.
- 3** Shut down node 1 and perform the following steps on the remaining node(s), one at a time:
 - a** Turn on the node.
 - b** Open **Disk Management**.
 - c** Assign the drive letters for the drives.
This procedure allows Windows to mount the volumes.
 - d** Reassign the drive letter, if necessary.
To reassign the drive letter:
 - With the mouse pointer on the same icon, right-click and select **Change Drive Letter and Path** from the submenu.
 - Click **Edit**, select the letter you want to assign the drive (for example, Z), and then click **OK**.
 - Click **Yes** to confirm the changes.
 - e** Power down the node.

If the cables are connected properly, the drive order is the same as is on each node, and the drive letter assignments of all the cluster nodes follow the same order as is on node 1. The volume labels can also be used to double-check the drive order by ensuring that the disk with volume label "Volume Z" is assigned to drive letter Z and so on for each disk on each node. Assign drive letters on each of the shared disks, even if the disk displays the drive letter correctly.

See your EMC documentation located on the Dell Support website at support.dell.com or the EMC support site located at www.emc.com for more information about the Navisphere Manager software.

Updating a Dell | EMC Storage System for Clustering

If you are updating an existing Dell | EMC storage system to meet the cluster requirements for the shared storage subsystem, you may need to install additional Fibre Channel disk drives in the shared storage system. The size and number of drives you add depend on the RAID level you want to use and the number of Fibre Channel disk drives currently in your system.

See your storage system's documentation for information on installing Fibre Channel disk drives in your storage system.

You may also need to upgrade the core software version that is running on the storage system or enable Access Logix. See the *Platform Guide* for specific version requirements.

Installing and Configuring MSCS

MSCS is a software component in Windows 2000, Advanced Server and an integrated service in Windows Server 2003. MSCS performs the basic cluster functionality, which includes membership, communication, and failover management. When MSCS is installed properly, the service starts on each node and responds automatically in the event that one of the nodes fails or goes offline. To provide application failover for the cluster, the MSCS software must be installed on each cluster node. See "Using MSCS" on page 87 for more information.

Installing and Configuring Microsoft Cluster Service (MSCS) with Windows 2000

After the Windows 2000 setup files are copied to the cluster node, Windows 2000 reboots and enters the second phase of installation.

During this phase, the Windows 2000 setup prompts you for the product key, licensing mode, system name, and additional information before starting the Windows 2000 Components program. From this dialog box, you can select the Windows 2000 core components for installation. Selecting **Cluster Service** copies the required files to the system.

If your operating system was preinstalled by Dell or you did not select Cluster Service when you installed the operating system, you can install the Cluster Service later by running **Add/Remove Components** from the **Control Panel**. After the Windows 2000 setup completes and all clustering prerequisites are met, run the Cluster Configuration Wizard to complete the installation, add additional node(s), and configure cluster resources.

To access the Cluster Configuration wizard:

- 1** Click the **Start** button and select **Run**.
- 2** In the **Run** box, type the following:
`C:\winnt\cluster\cluscfg.exe`
- 3** Click **OK**.
- 4** Follow the directions in the **Cluster Configuration Wizard** window.

Configuring Microsoft Cluster Service (MSCS) with Windows Server 2003

The cluster setup files are automatically installed on the system disk. To create a new cluster:

- 1** Click the **Start** button, select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 2** From the **File** menu, select **Open Connection**.
- 3** In the **Action** box of the **Open Connection to Cluster**, select **Create new cluster**.

The **New Server Cluster Wizard** window appears.

- 4** Click **Next** to continue.
- 5** Follow the procedures in the wizard, and then click **Finish**.
- 6** Add the additional node(s) to the cluster.
 - a** Turn on the remaining node(s).
 - b** Click the **Start** button, select **Programs**→**Administrative Tools**, and then double-click **Cluster Administrator**.
 - c** From the **File** menu, select **Open Connection**.
 - d** In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster**.

- e In the **Cluster or server name** box, type the name of the cluster or click **Browse** to select an available cluster from the list, and then click **OK**.

The **Add Nodes Wizard** window appears.

If the Add Nodes Wizard *does not* generate a cluster feasibility error, go to step f.

If the Add Nodes Wizard generates a cluster feasibility error, go to "Adding Cluster Nodes Using the Advanced Configuration Option" on page 76.

- f Click **Next** to continue.
- g Follow the procedures in the wizard and click **Finish**.

Adding Cluster Nodes Using the Advanced Configuration Option

If you are adding additional nodes to the cluster using the Add Nodes wizard and the nodes are not configured with identical internal storage devices, the wizard may generate one or more errors while checking cluster feasibility in the **Analyzing Configuration** menu. If this situation occurs, select **Advanced Configuration Option** in the Add Nodes wizard to add the nodes to the cluster.

To add the nodes using the **Advanced Configuration Option**:


- 1 From the **File** menu in Cluster Administrator, select **Open Connection**.
- 2 In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster** and then click **OK**.

The **Add Nodes Wizard** window appears.

- 3 Click **Next**.
- 4 In the **Select Computers** menu, click **Browse**.
- 5 In the **Enter the object names to select (examples)**, type the names of one to seven systems to add to the cluster, with each system name separated by a semicolon.
- 6 Click **Check Names**.

The Add Nodes Wizard verifies and underlines each valid system name.

- 7 Click **OK**.

- 8 In the **Select Computers** menu, click **Add**.
 - 9 In the **Advanced Configuration Options** window, click **Advanced (minimum) configuration**, and then click **OK**.
 - 10 In the **Add Nodes** window, click **Next**.
 - 11 In the **Analyzing Configuration** menu, Cluster Administrator analyzes the cluster configuration.
If Cluster Administrator discovers a problem with the cluster configuration, a warning icon appears in the **Checking cluster feasibility** window. Click the plus ("+") sign to review any warnings, if needed.
 - 12 Click **Next** to continue.
 - 13 In the **Password** field of the **Cluster Service Account** menu, type the password for the account used to run the Cluster Service, and click **Next**.
The **Proposed Cluster Configuration** menu appears with a summary with the configuration settings for your cluster.
 - 14 Click **Next** to continue.
The new systems (hosts) are added to the cluster. When completed, **Tasks completed** appears in the **Adding Nodes to the Cluster** menu.
-  **NOTE:** This process may take several minutes to complete.
- 15 Click **Next** to continue.
 - 16 In the **Completing the Add Nodes Wizard** window, click **Finish**.

Verifying Cluster Readiness

To ensure that your server and storage systems are ready for MSCS installation, ensure that these systems are functioning correctly and verify the following:

- All cluster servers are able to log on to the same domain.
- The shared disks are partitioned and formatted, and the same drive letters that reference logical drives on the shared storage system are used on each node.

All IP addresses and network names for each cluster node are communicating with each other and the public network

Installing Applications in the Cluster Group

The Cluster Group contains a network name and IP address resource, which is used to manage the cluster. Because the Cluster Group is dedicated to cluster management and for best cluster performance, Dell recommends that you do not install applications in this group.

Installing the Quorum Resource

When you install a Windows 2000 Advanced Server cluster, the installation wizard prompts you for a location to install the quorum resource. When you install a Windows Server 2003 cluster, the installation wizard automatically selects an NTFS disk as the quorum resource for you, which you can modify later. When you complete the procedures in the wizard, you can select another disk for the quorum using **Cluster Administrator**. To prevent quorum resource corruption, Dell recommends that you do not place applications or data on the disk.

Creating a LUN for the Quorum Resource

It is recommended that you create a separate LUN—approximately 1 GB in size—for the quorum resource.

When you create the LUN for the quorum resource:

- Format the LUN with NTFS.
- Use the LUN exclusively for your quorum logs.
- Do not store any application data or user data on the quorum resource.
- To easily identify the quorum resource, Dell recommends that you assign the drive letter "Q" to the quorum resource.



NOTE: The **Majority Node Set Quorum** types for Windows Server 2003 are not supported.

Preventing Quorum Resource Failure

Since the quorum resource plays a crucial role in cluster operation, losing a quorum resource causes the entire cluster to fail. To prevent cluster failure, configure the quorum resource on a RAID volume in the shared storage system.



NOTICE: It is recommended that you use a RAID level other than RAID 0, which is commonly called striping. RAID 0 configurations provide very high performance, but they do not provide the level of availability required for the quorum resource.

Configuring Windows 2000 Cluster Networks

When you install and configure a Windows 2000 Advanced Server cluster, the software installation wizard prompts you to identify the public and private network segments connected to your cluster nodes. Dell recommends the following configuration, which provides added fault tolerance for the private network:

- 1 Set the private network (cluster interconnect) to **Use for internal communications only**.
- 2 Name this network **Private**.
- 3 Set the client public network segment(s) to **All communications**.

This setting provides a redundant path for the cluster-to-cluster communication in the event the private network fails.

- 4 For all subsequent NICs, set each NIC to **Client use only or All communications**.
- 5 Set the priority of the networks so that the network you designated as **Private** has the highest priority for internal communications.

You can set the priority of the networks when you install MSCS or when using Cluster Administrator software.

Configuring Cluster Networks Running Windows Server 2003

When you install and configure a cluster running Windows Server 2003, the software installation wizard automatically configures all networks for mixed (public and private) use in your cluster. You can rename a network, allow or disallow the cluster to use a particular network, or modify the network role using **Cluster Administrator**. Dell recommends that you configure at least one network for the cluster interconnect (private network) and provide redundancy for the private network by configuring an additional network for mixed (public and private) use. If you have enabled network adapter teaming or are using dual-port NICs for use on your public network, you should change the configuration for these networks to support public communications only.

Verifying MSCS Operation

After you install MSCS, verify that the service is operating properly.

If you selected **Cluster Service** when you installed the operating system, see "Obtaining More Information" on page 80.

If you did not select **Cluster Service** when you installed the operating system:

- 1 Click the **Start** button and select **Programs**→**Administrative Tools**, and then select **Services**.
- 2 In the **Services** window, verify the following:
 - In the **Name** column, **Cluster Service** appears.
 - In the **Status** column, **Cluster Service** is set to **Started**.
 - In the **Startup Type** column, **Cluster Service** is set to **Automatic**.

Obtaining More Information

See Microsoft's online help for configuring the Cluster Service.

See "Using MSCS" on page 87 for more information on the Cluster Service.

Verifying Cluster Functionality

To verify cluster functionality, monitor the cluster network communications to ensure that your cluster components are communicating properly with each other. Also, verify that MSCS is running on the cluster nodes.

Verifying Cluster Resource Availability

In the context of clustering, a resource is a basic unit of failover management. Application programs are made up of resources that are grouped together for recovery purposes. All recovery groups, and therefore the resources that comprise the recovery groups, must be online (or in a ready state) for the cluster to function properly.

To verify that the cluster resources are online:

- 1** Start **Cluster Administrator** on the monitoring node.
- 2** Click the **Start** button and select **Programs**→ **Administrative Tools (Common)**→ **Cluster Administrator**.
- 3** Open a connection to the cluster and observe the running state of each resource group. If a group has failed, one or more of its resources might be offline.

Troubleshooting Failed Resources

Troubleshooting the failed resources is beyond the scope of this document, but examining the properties of each resource and ensuring that the specified parameters are correct are the first two steps in this process. In general, if a resource is offline, you can bring it online by right-clicking the resource and selecting **Bring Online** from the drop-down menu.

See the documentation and online help for Windows 2000 Advanced Server or Windows Server 2003, for information about troubleshooting resource failures.

Installing Your Cluster Management Software

This section provides information on configuring and administering your cluster using Microsoft® Cluster Administrator. Microsoft provides Cluster Administrator as a built-in tool for cluster management.

Microsoft Cluster Administrator

Cluster Administrator is Microsoft's tool for configuring and administering a cluster. The following procedures describe how to run Cluster Administrator locally on a cluster node and how to install the tool on a remote console.

Launching Cluster Administrator on a Cluster Node

- 1 Click the **Start** button and select **Programs**.
- 2 Select **Administrative Tools**.
- 3 Select **Cluster Administrator**.

Running Cluster Administrator on a Remote Console

You can administer and monitor the Cluster Service remotely by installing the Windows Administration Tools package and Cluster Administrator on a remote console (or management station) running the Microsoft Windows® operating system. Cluster Administrator is part of the Administration Tools package, which is included with the following operating systems:

- Windows 2000 Advanced Server
- Windows Server® 2003, Enterprise Edition
- Windows Server 2003, Enterprise x64 Edition

The Windows 2000 Administration Tools can only be installed on systems running Windows 2000. Additionally, the Windows 2003 Administrative Tools can only be installed on systems running Windows XP (with Service Pack 1 or later) and Windows Server 2003.

To install Cluster Administrator and the Windows Administration Tools package on a remote console:

- 1 Select a system that you wish to configure as the remote console.
- 2 Identify the operating system that is currently running on the selected system.
- 3 Insert the appropriate operating system media into the system's optical drive:
 - *Microsoft Windows 2000 Advanced Server CD*
 - *Windows Server 2003, Enterprise Edition CD*
 - *Windows Server 2003, Enterprise x64 Edition CD*
- 4 Open an Explorer window, navigate to the system's optical drive and double-click the `\i386` directory.
- 5 If you inserted the *Windows 2000 Advanced Server CD* or *Windows Server 2003, Enterprise Edition CD*, double-click **ADMINPAK.MSI**.
If you inserted the *Windows Server 2003, Enterprise x64 Edition CD*, double-click **WADMINPAK.MSI**.
- 6 Follow the instructions on your screen to complete the installation.

Launching Cluster Administrator on a Remote Console

Perform the following steps on the remote console:

- 1 Ensure that the Windows Administrative Tools package was installed on the system.
- 2 Click the **Start** button and select **Programs**.
- 3 Select **Administrative Tools**.
- 4 Select **Cluster Administrator**.

Installing Cluster Administrator for Windows Clusters on a Remote Console

You cannot install the Windows 2000 or Windows Server 2003 Administration Tools package on clients running any version of Windows NT[®] 4.0. However, a Windows 2000 Advanced Server or Windows Server 2003, Enterprise Edition cluster can be remotely administered using the Cluster Administrator included with Windows NT 4.0 operating systems with limited support.

See your Windows NT 4.0 operating system documentation for more information about the installation of Cluster Administrator on a remote client.

Cluster Administration and Monitoring

When using Cluster Administrator provided with Windows NT 4.0 operating systems on a system running Windows NT 4.0, Cluster Administrator may generate error messages if the software detects Windows 2000 or Windows Server 2003 cluster resources. It is strongly recommended to use client systems running Windows 2000 or Windows Server 2003 with the appropriate Administrator Pack for cluster administration and monitoring.

Using MSCS

Cluster Objects

Cluster objects are the physical and logical units managed by a cluster. Each object is associated with the following:

- Properties that define the object and its behavior within the cluster
- A set of cluster control codes used to manipulate the object's properties
- A set of object management functions to manage the object through MSCS

Cluster Networks

A cluster network provides a communications link between the cluster nodes (private network), the client systems in a local area network (public network), or a combination of the above (public-and-private network).

Preventing Network Failure

When you install MSCS, identify the public and private network segments connected to your cluster nodes. To ensure cluster failover and non-interrupted communications, perform the following procedures:

- Configure the private network for internal communications.
- Configure the public network for all communications to provide a redundant path if all of the private networks fail.
- Configure subsequent network adapters for client system use only or for all communications.

You can set priorities and roles of the networks when you install MSCS or when you use Microsoft® Cluster Administrator software.

Node-to-Node Communication

If a network is configured for public (client) access only, the Cluster Service will not use the network for internal node-to-node communications. If all of the networks that are configured for private (or mixed) communication fail, the nodes cannot exchange information and one or more nodes will terminate MSCS and temporarily stop participating in the cluster.

Network Interfaces

You can use Cluster Administrator or another cluster management application to view the state of all cluster network interfaces.

Cluster Nodes

A cluster node is a system in a cluster running the Microsoft® Windows® operating system and MSCS.

Each node in a cluster:

- Attaches to one or more cluster storage devices that store all of the cluster's configuration and resource data; nodes have access to all cluster configuration data
- Communicates with the other nodes through network adapters
- Is aware of systems that join or leave the cluster
- Is aware of the resources that are running on each node
- Is grouped with the remaining nodes under a common cluster name, which is used to access and manage the cluster

Table 5-1 defines states of a node during cluster operation

Table 5-1. Node States and Definitions

State	Definition
Down	The node is not actively participating in cluster operations.
Joining	The node is becoming an active participant in the cluster operations.
Paused	The node is actively participating in cluster operations but cannot take ownership of resource groups or bring resources online.

Table 5-1. Node States and Definitions (continued)

State	Definition
Up	The node is actively participating in all cluster operations, including hosting cluster groups.
Unknown	The node state cannot be determined.

When MSCS is configured on a node, the administrator chooses whether that node forms its own cluster or joins an existing cluster. When MSCS is started, the node searches for other active nodes on networks that are enabled for internal cluster communications.

Forming a New Cluster

MSCS maintains a current copy of the cluster database on all active nodes. If a node cannot join a cluster, the node attempts to gain control of the quorum resource and form a cluster. The node uses the recovery logs in the quorum resource to update its cluster database.

Joining an Existing Cluster

A node can join a cluster if it can communicate with another active node in the cluster. When a node joins a cluster, the node is updated with the latest copy of the cluster database. MSCS validates the node's name, verifies version compatibility, and the node joins the cluster.

Cluster Resources

A cluster resource is any physical or logical component that can be:

- Brought online and taken offline
- Managed in a cluster
- Hosted by one managed system at a time

When MSCS makes a resource request through a dynamic link library (DLL), the Resource Monitor checks and controls the resource's state.

Setting Resource Properties

Using the resource **Properties** dialog box, you can perform the following tasks:

- View or change the resource name, description, and possible owners.
- Assign a separate resource memory space.
- View the resource type, group ownership, and resource state.
- View which node currently owns the resource.
- View pre-existing dependencies and modify resource dependencies.
- Restart a resource and configure the resource settings (if required).
- Check the online state of the resource by configuring the **Looks Alive** (general check of the resource) and **Is Alive** (detailed check of the resource) polling intervals in MSCS.
- Specify the time requirement for resolving a resource in a pending state (**Online Pending** or **Offline Pending**) before MSCS places the resource in **Offline** or **Failed** status.
- Set specific resource parameters.

The **General**, **Dependencies**, and **Advanced** tabs are the same for every resource; however, some resource types support additional tabs.



NOTE: Do not update cluster object properties on multiple nodes simultaneously. See the MSCS online documentation for more information.

Resource Dependencies

MSCS uses the resource dependencies list when bringing resources online and offline. For example, if a group with a physical disk and a file share is brought online together, the physical disk containing the file share must be brought online before the file share. Table 5-2 shows resources and their dependencies.



NOTE: You must configure the required dependencies before you create the resource.

Table 5-2. Cluster Resources and Required Dependencies

Resource	Required Dependencies
File share	Network name (only if configured as a distributed file system [DFS] root)
IP address	None
Network name	IP address that corresponds to the network name
Physical disk	None

Setting Advanced Resource Properties

By using the **Advanced** tab in the **Properties** dialog box, you can perform the following tasks:

- Restart a resource or allow the resource to fail.
See "Adjusting the Threshold and Period Values" on page 93 for more information.
- Adjust the **Looks Alive** or **Is Alive** parameters.
- Select the default number for the resource type.
- Specify the time parameter for a resource in a pending state.

Resource Parameters

The **Parameters** tab in the **Properties** dialog box is available for most resources. Table 5-3 shows each resource and its configurable parameters.

Table 5-3. Resources and Configurable Parameters

Resource	Configurable Parameters
File share	Share permissions and number of simultaneous users Share name (clients systems detect the name in their browse or explore lists) Share comment Shared file path

Table 5-3. Resources and Configurable Parameters (continued)

Resource	Configurable Parameters
IP address	IP address Subnet mask Network parameters for the IP address resource (specify the correct network)
Network name	Cluster name or virtual server
Physical disk	Hard drive for the physical disk resource (cannot be changed after the resource is created)

Quorum Resource

Normally, the quorum resource is a common cluster resource that is accessible by all of the nodes. The quorum resource—typically a physical disk on a shared storage system—maintains data integrity, cluster unity, and cluster operations.

When the cluster is formed or when the nodes fail to communicate, the quorum resource guarantees that only one set of active communicating nodes is allowed to form a cluster. If a node fails and the node containing the quorum resource is unable to communicate with the remaining nodes, MSCS shuts down the node that does not control the quorum resource. If a node fails, the configuration database helps the cluster recover a failed resource or recreates the cluster in its current configuration.

The shared physical disk is the only resource supported by the solution that can act as a quorum resource.



NOTE: The Majority Node Set Quorum resource type is not supported.

Additionally, the quorum resource ensures cluster integrity. MSCS uses the quorum resource's recovery logs to update the private copy of the cluster database in each node, thereby maintaining the correct version of the cluster database and ensuring that the cluster is intact.

The operating system uses the quorum resource to ensure that only one set of active, communicating nodes is allowed to operate as a cluster. A node can form a cluster only if the node can gain control of the quorum resource. A node can join a cluster or remain in an existing cluster only if it can communicate with the node that controls the quorum resource.

Resource Failure

MSCS periodically launches the Resource Monitor to check if a resource is functioning properly. Configure the **Looks Alive** and **Is Alive** polls to check for failed resources. The **Is Alive** poll interval is typically longer than the **Looks Alive** poll interval because MSCS requests a more thorough check of the resource's state.



NOTE: Do not adjust the **Looks Alive** and **Is Alive** settings unless instructed to do so by technical support.

Adjusting the Threshold and Period Values

The **Threshold** value determines the number of attempts to restart the resource before the resource fails over. The **Period** value assigns a time requirement for the **Threshold** value to restart the resource.

If MSCS exceeds the maximum number of restart attempts within the specified time period and the failed resource has not been restarted, MSCS considers the resource to be failed.



NOTE: See "Setting Advanced Resource Properties" to configure the **Looks Alive**, **Is Alive**, **Threshold**, and **Period** values for a particular resource.



NOTE: Do not adjust the **Threshold** and **Period** settings unless instructed by technical support.

Configuring Failover

You can configure a resource to affect the group and fail over an entire group to another node when a resource fails in that group. If the number of failover attempts exceeds the group's threshold and the resource is still in a failed state, MSCS attempts to restart the resource after a period of time specified by the resource's **Retry Period On Failure** property.



NOTE: Do not adjust the **Retry Period On Failure** settings unless instructed by technical support.

When you configure **Retry Period On Failure**, use the following guidelines:

- Select a unit value of minutes rather than milliseconds (the default value is milliseconds).
- Select a value that is greater than or equal to the value of the resource's restart period property.

Resource Dependencies

A dependent resource requires another resource to operate. Table 5-4 describes resource dependencies.

Table 5-4. Resource Dependencies

Term	Definition
Dependent resource	A resource that depends on other resources.
Dependency	A resource on which another resource depends.
Dependency tree	A series of dependency relationships or hierarchy. The following rules apply to a dependency tree: <ul style="list-style-type: none">• A dependent resource and its dependencies must be in the same group.• A dependent resource is taken offline before its dependencies and brought online after its dependencies, as determined by the dependency hierarchy.

Creating a New Resource

Before you add a resource to your cluster solution, verify that the following conditions exist in your cluster:

- The type of resource is either a standard resource provided with MSCS or a custom resource provided by Microsoft or a third party vendor.
- A group that will contain the resource already exists within your cluster.
- All dependent resources have been created.
- A separate Resource Monitor exists (recommended for any resource that has caused problems in the past).

To create a new resource:

- 1 Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.

The **Cluster Administrator** window appears.

- 2 In the console tree, double-click the **Groups** folder.
- 3 Select the group to which you want the resource to belong.

- 4 On the **File** menu, point to **New** and click **Resource**.
- 5 In the New Resource wizard, type the appropriate information in the **Name** and **Description** fields and select the appropriate **Resource type** and **Group** for the new resource.
- 6 Click **Next**.
- 7 Add or remove possible owners of the resource and click **Next**.
The **New Resource** window appears with **Available resources** and **Resource dependencies** selections.
 - To add dependencies, under **Available resources**, select a resource, and then click **Add**.
 - To remove dependencies, under **Resource dependencies**, select a resource, and then click **Remove**.
- 8 Repeat step 7 for all resource dependencies and click **Finish**.
- 9 Set the resource properties.
For more information about setting resource properties, see the MSCS online help.

Deleting a Resource

- 1 Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
The **Cluster Administrator** window appears.
- 2 In the console tree, double-click the **Resources** folder.
- 3 In the details pane, select the resource that you want to remove.
- 4 In the **File** menu, click **Offline**.
The resource must be taken offline before it can be deleted.
- 5 In the **File** menu, click **Delete**.

When you delete a resource, Cluster Administrator deletes all of the resources that are dependent on the deleted resource.

File Share Resource Type

If you want to use your cluster solution as a high-availability file server, select one of the following types of file share for your resource:

- Basic file share — Publishes a file folder to the network under a single name.
- Share subdirectories — Publishes several network names—one for each file folder and all of its immediate subfolders. This method is an efficient way to create large numbers of related file shares on a file server.
- Distributed File System (DFS) root — Creates a resource that manages a stand-alone DFS root. Fault-tolerant DFS roots cannot be managed by this resource. A DFS root file share resource has required dependencies on a network name and an IP address. The network name can be either the cluster name or any other network name for a virtual server.

Configuring Active and Passive Cluster Nodes

Active nodes process application requests and provide client services. Passive nodes are backup nodes that ensure that client applications and services are available if a hardware or software failure occurs. Cluster configurations may include both active and passive nodes.



NOTE: Passive nodes must be configured with appropriate processing power and storage capacity to support the resources that are running on the active nodes.

Your cluster solution supports variations of active/active (active^x) and active/passive (active^x/passive^x) configurations. The variable *x* indicates the number of nodes that are active or passive.

Cluster solutions running the Windows Server 2003 operating system can support up to eight nodes in multiple configurations as shown in Table 5-6. Cluster solutions running Windows 2000 Advanced Server can support only two nodes and are limited to active² and active¹/passive¹ configurations because this solution supports only two nodes.

An active/active (active^x) configuration contains virtual servers running separate applications or services on each node. When an application is running on node 1, the remaining node(s) do not have to wait for node 1 to fail. Those node(s) can run their own cluster-aware applications (or another instance of the same application) while providing failover for the resources on node 1. For example, multiway failover is an active/active failover solution

because running applications from a failed node can migrate to multiple active nodes in the cluster. However, you must ensure that adequate resources are available on each node to handle the increased load if one node fails.

In an active/passive (active^x/passive^x) configuration, one or more *active* cluster nodes are processing requests for a clustered application while the *passive* cluster nodes only wait for the active node(s) to fail.

Table 5-5 provides a description of active/active configuration types.

Table 5-5. Active/Active Configuration Types

Configuration Type	Active Cluster Node(s)	Definition
Active ²	2	<p>The active node(s) processes requests and provides failover for each other, depending on node resources and your configuration.</p> <p>NOTE: Configurations running Windows 2000 Advanced Server are limited to two-node, active/active (active²) configurations.</p>
Active ³	3	
Active ⁴	4	
Active ⁵	5	
Active ⁶	6	
Active ⁷	7	
Active ⁸	8	

Table 5-6 provides a description of some active/passive configuration types.

Table 5-6. Active/Passive Configuration Types

Configuration Type	Active Cluster Node(s)	Passive Cluster Node(s)	Description
Active ¹ /Passive ¹	1	1	The active node(s) processes requests while the passive node waits for the active node to fail.
Active ² /Passive ¹	2	1	
Active ² /Passive ²	2	2	NOTE: Configurations running Windows 2000 Advanced Server are limited to two-node, active/passive (active ¹ /passive ¹) configurations.
Active ³ /Passive ¹	3	1	
Active ³ /Passive ²	3	2	
Active ⁴ /Passive ¹	4	1	
Active ⁴ /Passive ²	4	2	
Active ⁵ /Passive ¹	5	1	
Active ⁵ /Passive ²	5	2	
Active ⁶ /Passive ¹	6	1	
Active ⁶ /Passive ²	6	2	
Active ⁷ /Passive ¹	7	1	

Failover Policies

When implementing a failover policy, configure failback if the cluster node lacks the resources (such as memory or processing power) to support cluster node failures.

Windows 2000 Advanced Server Cluster Configurations

Configurations running Windows 2000 Advanced Server include only two nodes. Therefore, failover to one node is the only option.

Windows Server 2003 Cluster Configurations

Cluster configurations running Windows Server 2003 provide the following failover policies:

- N (number of active nodes) + I (number of inactive nodes) failover
- Failover pair
- Multiway failover
- Failover ring

Table 5-7 provides an overview of the failover policies implemented with Windows 2000 Advanced Server and Windows Server 2003. For more information, see the sections that follow this table.

Table 5-7. Windows Server 2003 Failover Policies

Failover Policy	Description	Advantage	Disadvantage(s)
$N + I$	One or more nodes provides backup for multiple servers.	Highest resource availability.	<ul style="list-style-type: none">• May not handle more than one backup node failure.• May not fully utilize all of the nodes.
Failover pair	Applications can fail over between the two nodes.	Easy to plan the capacity of each node.	Applications on the pair cannot tolerate two node failures.

Table 5-7. Windows Server 2003 Failover Policies (continued)

Failover Policy	Description	Advantage	Disadvantage(s)
Multiway	Running applications migrate to multiple nodes in the cluster.	Application load balancing.	Must ensure that the failover nodes have ample resources available to handle the additional workload.
Failover ring	Running applications migrate to the next preassigned node.	Easy to scope node capacity for one server failure.	The next node for failover may not have ample resources to handle the workload.

N + I Failover

N + I failover is an active/passive policy where dedicated passive cluster node(s) provide backup for the active cluster node(s). This solution is best for critical applications that require dedicated resources. However, backup nodes add a higher cost of ownership because they remain idle and do not provide the cluster with additional network resources.

Figure 5-1 shows an example of a 6 + 2 (*N + I*) failover configuration with six active nodes and two passive nodes. Table 5-8 provides an *N + I* failover matrix for Figure 5-1.

Figure 5-1. Example of an N+I Failover Configuration for an Eight-Node Cluster




Table 5-8. Example of an N+I Failover Configuration for an Eight-Node Cluster

Cluster Resource Group	Primary Node	AntiAffinityClassNames Value
A	Node 1	AString
B	Node 2	AString
C	Node 3	AString
D	Node 4	AString
E	Node 5	AString
F	Node 6	AString

Configuring Group Affinity

On $N + I$ (active/passive) failover clusters running Windows Server 2003, some resource groups may conflict with other groups if they are running on the same node. For example, running more than one Microsoft Exchange virtual server on the same node may generate application conflicts. Use Windows Server 2003 to assign a public property (or attribute) to a dependency between groups to ensure that they fail over to similar or separate nodes. This property is called *group affinity*.

Group affinity uses the `AntiAffinityClassNames` public property, which ensures that designated resources are running on *separate nodes*, if possible.

For example, in Table 5-8, the `AntiAffinityClassNames` string for cluster resource group A and group B are identical (`AString`), which indicates that these groups are assigned to run on separate nodes, if possible. If node 1 fails, resource group A will fail over to the next backup node (node 7). If node 2 then fails, because their `AntiAffinityClassNames` string value (`AString`) identifies group A and group B as conflicting groups, group B will skip node 7 and instead fail over to node 8.

To set the public property for the cluster groups shown in Table 5-8:

- 1 Open a command prompt.
- 2 Type the following:

```
cluster group "A" /prop AntiAffinityClassNames=  
"AString"
```

- 3 Repeat step 2 for the remaining cluster groups.

Use the "Cluster Data Form" on page 129 to specify group affinity in your $N + I$ cluster configuration.

Failover Pair

Failover pair is a policy in which each application can fail over between two specific nodes in a multinode cluster. The **Possible Owners** list in Cluster Administrator determines which nodes run the failed over applications.

If you have applications that run well on two-node, active/active configurations running Windows 2000 Advanced Server, and you want to migrate these applications to Windows Server 2003, failover pair is a good policy. This solution is easy to plan and administer, and applications that do

not run well on the same server can easily be moved into separate failover pairs. However, in a failover pair, applications on the pair cannot tolerate two node failures.

Figure 5-2 shows an example of a failover pair configuration. Table 5-9 provides a failover configuration for the cluster shown in Figure 5-2.

Figure 5-2. Example of a Failover Pair Configuration




Table 5-9. Example of a Failover Pair Configuration for a Four-Node Cluster

Cluster Resource Group	Possible Owners List
App1	1, 2
App2	3, 4

Multiway Failover

Multiway failover is an active/active policy where running applications from a failed node migrate to multiple nodes in the cluster. This solution provides automatic failover and load-balancing. Ensure that the failover nodes have sufficient resources to handle the workload. Figure 5-3 shows an example of four-node multiway failover configuration.

Table 5-10 shows a four-node multiway failover configuration for the cluster shown in Figure 5-3. For each resource group, the failover order in the **Preferred Owners** list in Cluster Administrator outlines the order that you want that resource group to fail over. In this example, node 1 owns applications A, B, and C. If node 1 fails, applications A, B, and C fail over to cluster nodes 2, 3, and 4. Configure the applications similarly on nodes 2, 3, and 4.

When implementing multiway failover, configure failback to avoid performance degradation. See "Using MSCS" on page 87 for more information.

Figure 5-3. Example of a Four-Node Multiway Failover Configuration




Table 5-10. Example of a Four-Node Multiway Failover Configuration


Application	Failover Order in the Preferred Owners List
A	Node 2
B	Node 3
C	Node 4

Failover Ring

Failover ring is an active/active policy where all running applications migrate from the failed node to the next preassigned node in the Preferred Owners List. If the failing node is the last node in the list, the failed node's applications fail over to the first node.

While this type of failover provides high availability, ensure that the next node for failover has sufficient resources to handle the additional workload. Figure 5-4 shows an example of a failover ring configuration.

Figure 5-4. Example of a Four-Node Failover Ring Configuration



Failover and Failback Capabilities

Failover

When an application or cluster resource fails, MSCS detects the failure and attempts to restart the resource. If the restart fails, MSCS takes the application offline, moves the application and its resources to another node, and restarts the application on the other node.

See "Setting Advanced Resource Properties" for more information.

Cluster resources are placed in a group so that MSCS can move the resources as a combined unit, ensuring that the failover and/or failback procedures transfer all resources.

After failover, Cluster Administrator resets the following recovery policies:

- Application dependencies
- Application restart on the same node
- Workload rebalancing (or failback) when a failed node is repaired and brought back online

Failback

Failback returns the resources back to their original node. When the system administrator repairs and restarts the failed node, MSCS takes the running application and its resources offline, moves them from the failover cluster node to the original node, and then restarts the application.

You can configure failback to occur immediately, at any given time, or not at all. To minimize the delay until the resources come back online, configure the failback time during off-peak hours.

Modifying Your Failover Policy

Use the following guidelines when you modify your failover policy:

- Define how MSCS detects and responds to group resource failures.
- Establish dependency relationships between the resources to control the order in which the resources are taken offline.
- Specify time-out, failover threshold, and failover period for your cluster resources.
See "Setting Advanced Resource Properties" for more information.
- Specify a Possible Owner List in Microsoft Cluster Administrator for cluster resources. The Possible Owner List for a resource controls which nodes are allowed to host the resource.

See the Cluster Administrator documentation for more information.

Upgrading to a Cluster Configuration

Before You Begin

Before you upgrade your non-clustered system to a cluster solution:

- Back up your data.
- Verify that your hardware and storage systems meet the minimum system requirements for a cluster as described in "System Requirements" on page 18.
- Verify that your hardware and storage systems are installed and configured as explained in the following sections:
 - "Cabling Your Cluster Hardware" on page 23
 - "Preparing Your Systems for Clustering" on page 47
 - "Installing Your Cluster Management Software" on page 83

Supported Cluster Configurations

Dell certifies and supports only solutions that are configured with the Dell products described in this guide. For a description of the cluster components, see the *Platform Guide*.

Completing the Upgrade

After installing the required hardware and network adapter upgrades, set up and cable the system hardware.



NOTE: You may need to reconfigure your switch or storage groups so that both nodes in the cluster can access their LUNs.

The final phase for upgrading to a cluster solution is to install and configure Windows 2000 Advanced Server or Windows Server 2003 with MSCS.

Upgrading Your Operating System

You can upgrade your Windows® 2000 Advanced Server cluster to a Windows Server 2003 cluster using one of the following methods:

- Standard upgrade — Upgrade the operating system on each cluster node while all cluster nodes are disconnected from the client network. This procedure requires you to re-create your cluster configuration.
See your Windows operating system documentation for performing a standard upgrade.
- Rolling upgrade — Upgrade the operating system on each cluster node while the remaining cluster nodes are connected to the client network and available to handle client requests. This procedure *does not* require you to re-create your cluster configuration. However, each cluster node must be configured with the appropriate resources to run all virtual servers and services for the entire cluster while you upgrade the remaining node.



NOTE: You cannot upgrade to Windows Server 2003, Enterprise x64 Edition. Only a new installation is permitted for Windows Server 2003, Enterprise x64 Edition.

The following section explains how to perform a rolling upgrade on a two-node cluster.

Performing a Rolling Upgrade

Before you perform a rolling upgrade:


- Ensure that your cluster nodes are running Windows 2000 Advanced Server.
- Back up your data and system states.
- Run the Check System Compatibility Wizard to determine if your cluster nodes are configured with the appropriate resources to run Windows Server 2003, Enterprise Edition.
- Ensure that your cluster service account is an explicit member of the local administrator's group and that its user-right privileges are set to **Act as part of operating system**.



NOTE: Some cluster resources do not support a rolling upgrade. For more information, see:

<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/cluster/rlupnet.mspx>

Upgrading Node 1


- 1** Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
- 2** In Cluster Administrator, right-click a node and then click **Pause Node**.
The following steps refer to the node you selected as node 1.
- 3** Right-click a cluster group and then click **Move Group**.
The cluster group is moved and restarted on node 2.
- 4** Repeat step 3 for the remaining cluster groups.
- 5** Uninstall the EMC® PowerPath® software.
See your EMC PowerPath software documentation for more information.
- 6** Insert the *Microsoft Windows Server 2003, Enterprise Edition* CD into the CD drive.
- 7** Double-click **Install Windows Server 2003, Enterprise Edition**.
The **Windows Setup Wizard** window appears.
- 8** Follow the instructions in the Windows Setup Wizard to upgrade your operating system.
 **NOTE:** If you are running IIS World Wide Web Publishing Service on your cluster node, this service is disabled during the upgrade to protect your system.
- 9** Reinstall the HBA drivers, PowerPath software, and any other required software applications.
See your EMC PowerPath software documentation for more information.
- 10** Verify that the upgraded node is functioning correctly.
 - a** Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
 - b** Move one or more cluster resource groups from node 2.
 - c** Verify that the resource group(s) can be brought online.
 - d** Close Cluster Administrator.
 - e** Remove the CD from the CD drive.

11 Go to "Upgrading Node 2" on page 110.



NOTE: After you upgrade node 1, your cluster is running two separate operating systems. It is recommended that you do not modify your cluster configuration—such as adding or removing cluster nodes or resources—until you upgrade both cluster nodes.

Upgrading Node 2

- 1 On node 2, click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
- 2 In Cluster Administrator, right-click node 1 and then click **Resume Node**.
- 3 Right-click node 2 and then click **Pause Node**.
- 4 Right-click a cluster group and then click **Move Group**.
The cluster group is moved and restarted on node 1.
- 5 Repeat step 4 for the remaining cluster groups.
- 6 Uninstall the EMC PowerPath software.
See your EMC PowerPath software documentation for more information.
- 7 Insert the *Microsoft Windows Server 2003, Enterprise Edition* CD into the CD drive.
- 8 Double-click **Install Windows Server 2003, Enterprise Edition**.
The **Windows Setup Wizard** window appears.
- 9 Follow the instructions in the Windows Setup Wizard to upgrade your operating system.
 **NOTE:** If you are running IIS World Wide Web Publishing Service on your cluster node, this service is disabled during the upgrade to protect your system.
- 10 Reinstall the HBA drivers, PowerPath software, and any other required software applications.
See your EMC PowerPath software documentation for more information.

- 11** Verify that the upgraded node is functioning correctly.
 - a** Click the **Start** button and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
 - b** Move one or more cluster resource groups from node 1.
 - c** Verify that the resource group(s) can be brought online.
 - d** Remove the CD from the CD drive.
- 12** In Cluster Administrator, redistribute the cluster groups to the appropriate cluster nodes.
- 13** Close Cluster Administrator.

Maintaining Your Cluster

Adding a Network Adapter to a Cluster Node



NOTE: To perform this procedure, the following must be installed on both nodes: Microsoft Windows 2000 Advanced Server or Windows Server 2003 (including the latest service packs) and MSCS.

- 1 Move all resources from the node you are upgrading to another node. See the MSCS documentation for information about moving cluster resources to a specific node.
- 2 Shut down the node you are upgrading.
- 3 Install the additional network adapter.
See the system's *Installation and Troubleshooting Guide* for expansion card installation instructions.
- 4 Turn on the node and allow the Windows operating system to boot. Windows detects the new adapter and installs the appropriate drivers.



NOTE: If Windows *does not* detect the new network adapter, the network adapter is not supported.

- 5 Update the network adapter drivers (if required).
- 6 Configure the network adapter addresses:
 - a Click the **Start** button, select **Control Panel**, and then double-click **Network Connections**.
 - b In the **Connections** box, locate the new adapter that you installed in the system.
 - c Right-click the new adapter and select **Properties**.
 - d Assign a unique static IP address, subnet mask, and gateway.



NOTE: Ensure that the host ID portion of the new network adapter's IP address is different from that of the first network adapter. For example, if the first network adapter in the node had an address of 192.168.1.101 with a subnet mask of 255.255.255.0, for the second network adapter you might assign the IP address 192.168.2.102 and the subnet mask 255.255.255.0.

- 7 Click **OK** and exit the network adapter properties.
- 8 Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
- 9 Click the **Network** tab.
- 10 Verify that a new resource labeled "New Cluster Network" appears in the window.
To rename the new resource, right-click the resource and enter a new name.
- 11 Move all cluster resources back to the original node.
- 12 Repeat step 2 through step 11 on each node.



NOTE: For each node, ensure that you assign the IP address on the same subnet as you did on the first node.

If the installation and IP address assignments have been performed correctly, all of the new network adapter resources appear online and respond successfully to ping commands.

Changing the IP Address of a Cluster Node on the Same IP Subnet



NOTE: If you are migrating your nodes to a different subnet, take all cluster resources offline and then migrate all nodes together to the new subnet.

- 1 Open **Cluster Administrator**.
- 2 Stop **MSCS** on the node.
The Cluster Administrator utility running on the second node indicates that the first node is down by displaying a red icon in the **Cluster Service** window.
- 3 Reassign the IP address.
- 4 If you are running **DNS**, verify that the **DNS** entries are correct (if required).
- 5 Restart **MSCS** on the node.

The nodes re-establish their connection and Cluster Administrator changes the node icon back to blue to show that the node is back online.

Uninstalling MSCS From Clusters Running Microsoft Windows 2000 Advanced Server

- 1 Take all resource groups offline or move them to another cluster node.
- 2 Stop MSCS on the node that you want to uninstall.
- 3 Click the **Start** button and select **Settings**→ **Control Panel**→ **Add/Remove Programs**.
- 4 Select **Add/Remove Windows Components**.
- 5 Deselect the check box for MSCS, click **Next**, and then click **Finish**.
- 6 From the remaining node, click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
- 7 Right-click the node icon and select **Options**→ **Evict Node**.
- 8 Close Cluster Administrator.


Removing Nodes From Clusters Running Microsoft Windows Server 2003

- 1 Move all resource groups to another cluster node.
- 2 Click the **Start** button, select **Programs**→ **Administrative Tools**, and then double-click **Cluster Administrator**.
- 3 In Cluster Administrator, right-click the icon of the node you want to uninstall and then select **Stop Cluster Service**.
- 4 In Cluster Administrator, right-click the icon of the node you want to uninstall and then select **Evict Node**.

If you cannot evict the node or the node is the last node in the cluster:

- a Open a command prompt.
 - b Type `cluster node <node_name> /force`
where `<node_name>` is the cluster node you are evicting from the cluster.
- 5 Close Cluster Administrator.

Running `chkdsk /f` on a Quorum Resource

 **NOTE:** You cannot run the `chkdsk` command with the `/f` (fix) option on a device that has an open file handle active. Because MSCS maintains an open file handle on the quorum resource, you cannot run `chkdsk /f` on the hard drive that contains the quorum resource.

- 1 Move the quorum resource temporarily to another drive.
- 2 Right-click the cluster name and select **Properties**.
- 3 Click the **Quorum** tab.
- 4 Select another disk as the quorum resource and press <Enter>.
- 5 Run `chkdsk /f` on the drive that previously stored the quorum resource.
- 6 Move the quorum resource back to the original drive.

Recovering From a Corrupt Quorum Disk

The quorum disk maintains the configuration data necessary for recovery when a node fails. If the quorum disk resource is unable to come online, the cluster does not start and all of the shared drives are unavailable. If this situation occurs and you must run `chkdsk` on the quorum disk, start the cluster manually from the command line.

To start the cluster manually from a command line prompt:

- 1 Open a command line window.
- 2 Select the cluster directory by typing one of the following:
 - `cd \winnt\cluster` (for Microsoft Windows 2000 Advanced Server)
 - Or
 - `cd \windows\cluster` (for Windows Server 2003)
- 3 Start MSCS in manual mode (on one node only) with no quorum logging by typing the following:
`clussvc -debug -noquorumlogging`
MSCS starts.

- 4 Run `chkdsk /f` on the disk designated as the quorum resource:
 - a Open a second command line window.
 - b Type `chkdsk /f`.
- 5 After the `chkdsk` utility completes, stop MSCS by pressing `<Ctrl><c>` in the first command line window.
- 6 Restart MSCS from the Services console:
 - a Click the **Start** button and select **Programs**→**Administrative Tools**→**Services**.
 - b In the **Services** window, right-click **Cluster Service**.
 - c In the drop-down menu, click **Start**.
 - d At the command line prompt in either window, type `Net Start Clussvc`.
The Cluster Service restarts.

See the Microsoft Knowledge Base article 258078 located on the Microsoft Support website at www.microsoft.com for more information on recovering from a corrupt quorum disk.

Changing the MSCS Account Password in Windows Server 2003

To change the service account password for all nodes running Microsoft Windows Server 2003, type the following at a command line prompt:

```
Cluster /cluster:[cluster_name] /changepass
```

where *cluster_name* is the name of your cluster


For help changing the password, type:

```
cluster /changepass /help
```



NOTE: Windows Server 2003, does not accept blank passwords for MSCS accounts.

Reformatting a Cluster Disk

 **NOTICE:** Ensure that all client systems are disconnected from the cluster disk before you perform this procedure.

- 1 Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
- 2 In the **Cluster Administrator** left pane, expand the **Groups** directory.
- 3 In the **Groups** directory, right-click the cluster resource group that contains the disk to be reformatted and select **Take Offline**.
- 4 In the **Cluster Administrator** right pane, right-click the physical disk you are reformatting and select **Bring Online**.
- 5 In the **Cluster Administrator** right pane, right-click the physical disk you are reformatting and select **Properties**.

The **Properties** window appears.

- 6 Click the **Advanced** tab.
- 7 In the "Looks Alive" **poll interval** box, select **Specify value**.
- 8 In the **Specify value** field, type:

6000000

where 6000000 equals 6,000,000 milliseconds (100 minutes)

- 9 Click **Apply**.
- 10 On the Windows desktop, right-click the **My Computer** icon and select **Manage**.

The **Computer Management** window appears.

- 11 In the **Computer Management** left pane, click **Disk Management**.
The physical disk information appears in the right pane.

- 12 Right-click the disk you want to reformat and select **Format**.
Disk Management reformats the disk.

- 13** In the **File** menu, select **Exit**.
- 14** In the "**Looks Alive**" **poll interval** box, select **Use value from resource type** and click **OK**.
- 15** In the **Cluster Administrator** left pane, right-click the cluster group that contains the reformatted disk and select **Bring Online**.
- 16** In the **File** menu, select **Exit**.

A

Troubleshooting

This appendix provides troubleshooting information for your cluster configuration.

Table A-1 describes general cluster problems you may encounter and the probable causes and solutions for each problem.

Table A-1. General Cluster Troubleshooting

Problem	Probable Cause	Corrective Action
The nodes cannot access the storage system, or the cluster software is not functioning with the storage system.	The storage system is not cabled properly to the nodes or the cabling between the storage components is incorrect.	Ensure that the cables are connected properly from the node to the storage system. See "Cabling Your Cluster For Public and Private Networks" on page 26 for more information.
	The length of the interface cables exceeds the maximum allowable length.	Ensure that the fiber optic cables do not exceed 300 m (multimode) or 10 km (single mode switch-to-switch connections only).
	One of the cables is faulty.	Replace the faulty cable.
Access Control is not enabled correctly.		Verify the following: <ul style="list-style-type: none">• All switched zones are configured correctly.• The EMC® Access Logix™ software is enabled on the storage system.• All LUNs and hosts are assigned to the proper storage groups.
	The cluster is in a SAN, and one or more zones are not configured correctly.	Verify the following: <ul style="list-style-type: none">• Each zone contains only one initiator (Fibre Channel daughter card).• Each zone contains the correct initiator and the correct storage port(s).

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
One of the nodes takes a long time to join the cluster . OR One of the nodes fail to join the cluster.	The node-to-node network has failed due to a cabling or hardware failure. Long delays in node-to-node communications may be normal.	Check the network cabling. Ensure that the node-to-node interconnection and the public network are connected to the correct NICs. Verify that the nodes can communicate with each other by running the ping command from each node to the other node. Try both the host name and IP address when using the ping command.
	One or more nodes may have the Internet Connection Firewall enabled, blocking Remote Procedure Call (RPC) communications between the nodes.	Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster Service (MSCS) and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Attempts to connect to a cluster using Cluster Administrator fail.	The Cluster Service has not been started.	Verify that the Cluster Service is running and that a cluster has been formed. Use the Event Viewer and look for the following events logged by the Cluster Service:
	A cluster has not been formed on the system. The system has just been booted and services are still starting.	<p>Microsoft Cluster Service successfully formed a cluster on this node.</p> <p>or</p> <p>Microsoft Cluster Service successfully joined the cluster.</p> <p>If these events do not appear in Event Viewer, see the Microsoft® Cluster Service Administrator's Guide for instructions on setting up the cluster on your system and starting the Cluster Service.</p>
	The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes.	<p>Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster Service (MSCS) and the clustered applications or services.</p> <p>See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.</p>

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
You are prompted to configure one network instead of two during MSCS installation.	The TCP/IP configuration is incorrect.	The node-to-node network and public network must be assigned static IP addresses on different subnets. See "Assigning Static IP Addresses to Cluster Resources and Components" on page 52 for information about assigning the network IPs.
	The private (point-to-point) network is disconnected.	Ensure that all systems are powered on so that the NICs in the private network are available.
Using Microsoft Windows NT® 4.0 to remotely administer a Windows® 2000 Advanced Server or Windows Server 2003 cluster generates error messages.	Normal. Some resources in Windows 2000, Advanced Server and Windows Server 2003 are not supported in Windows NT 4.0.	Dell strongly recommends that you use Windows 2000 Professional, Server, or Advanced server for remote administration of a cluster running Windows 2000 Advanced Server. Similarly, use Windows XP Professional or Windows Server 2003 for remote administration of a cluster running Windows Server 2003.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Unable to add a node to the cluster.	<p>The new node cannot access the shared disks.</p> <p>The shared disks are enumerated by the operating system differently on the cluster nodes.</p>	<p>Ensure that the new cluster node can enumerate the cluster disks using Windows Disk Administration. If the disks do not appear in Disk Administration, check the following:</p> <ul style="list-style-type: none">• Check all cable connections• Check all zone configurations• Check the Access Control settings on the attached storage systems• Use the "Advanced" with "Minimum" option
	<p>One or more nodes may have the Internet Connection Firewall enabled, blocking Remote Procedure Call (RPC) communications between the nodes.</p>	<p>Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster Service (MSCS) and the clustered applications or services.</p> <p>See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.</p>
The disks on the shared cluster storage appear unreadable or uninitialized in Windows Disk Administration	<p>This situation is normal if you stopped the Cluster Service. If you are running Windows Server 2003, this situation is normal if the cluster node does not own the cluster disk.</p>	<p>No action required.</p>

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Cluster Services does not operate correctly on a cluster running Windows Server 2003, Enterprise Edition with Service Pack 1 (SP1) or Windows Server 2003, Enterprise x64 Edition and the Internet Firewall enabled.	The Windows Internet Connection Firewall is enabled, which may conflict with Cluster Services.	<p>Perform the following steps:</p> <ol style="list-style-type: none">1 On the Windows desktop, right-click My Computer and click Manage.2 In the Computer Management window, double-click Services.3 In the Services window, double-click Cluster Services.4 In the Cluster Services window, click the Recovery tab.5 Click the First Failure drop-down arrow and select Restart the Service.6 Click the Second Failure drop-down arrow and select Restart the service.7 Click OK. <p>For information on how to configure your cluster with the Windows Internet Connection Firewall enabled, see Microsoft Base (KB) articles 258469 and 883398 at the Microsoft Support website at support.microsoft.com and the Microsoft Windows Server 2003 Technet website at www.microsoft.com/technet.</p>

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Public network clients cannot access the applications or services that are provided by the cluster.	One or more nodes may have the Internet Connection Firewall enabled, blocking Remote Procedure Call (RPC) communications between the nodes.	Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster Service (MSCS) and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.

Cluster Data Form

You can attach the following form in a convenient location near each cluster node or rack to record information about the cluster. Use the form when you call for technical support.

Table B-1.

Cluster Information	Cluster Solution
Cluster name and IP address	
Server type	
Installer	
Date installed	
Applications	
Location	
Notes	

Table B-2.

Node Name	Service Tag Number	Public IP Address	Private IP Address

Additional Networks

Table B-3.

Array	Array xPE Type	Array Service Tag Number or World Wide Name Seed	Number of Attached DAEs
1			
2			
3			
4			

Zoning Configuration Form

Node	HBA WWPNS or Alias Names	Storage WWPNS or Alias Names	Zone Name	Zone Set for Configuration Name

Index

A

- Access Control
 - about, 62
- Access Logix
 - about, 61
- active/active
 - about, 96

C

- cable configurations
 - cluster interconnect, 28
 - for client networks, 27
 - for mouse, keyboard, and monitor, 23
 - for power supplies, 23
- chkdsk/f
 - running, 116
- cluster
 - cluster objects, 87
 - forming a new cluster, 89
 - joining an existing cluster, 89
 - optional configurations, 17
 - verifying functionality, 80
 - verifying readiness, 77
 - verifying resource availability, 80
- Cluster Administrator
 - about, 83
- cluster configurations
 - active/active, 96
 - active/passive, 96
 - connecting to multiple shared storage systems, 43
 - connecting to one shared storage system, 17
 - direct-attached, 17, 31
 - SAN-attached, 18
 - supported configurations, 107
- cluster group
 - installing applications, 78
- cluster networks
 - configuring Windows 2000 cluster networks, 79
 - configuring Windows Server 2003 cluster networks, 79
- cluster nodes
 - about, 88
 - states and definitions, 88
- cluster objects
 - about, 87
- cluster resources
 - configurable parameters, 91
 - resource dependencies, 94
 - resource failure, 93
 - setting resource properties, 90
- Cluster Service
 - uninstalling, 115

- cluster storage
 - requirements, 20

- clustering
 - overview, 9

- connectors
 - about, 29
 - applications, 30
 - duplex LC multimode, 29
 - duplex SC multimode, 29

D

- Dell | EMC CX600
 - cabling the cluster nodes, 31
 - cabling the cluster nodes in a SAN-attached environment, 37
 - cabling to one cluster, 31
 - cabling to one SAN-attached cluster, 37
 - cabling to two clusters, 35
 - configuring, 60
 - installing, 60
 - updating for cluster use, 74
 - zoning in a switched environment, 43
- direct-attached cluster
 - about, 31
 - cabling, 31
- domain model
 - selecting, 49

- drive letters
 - assigning to shared storage systems, 69

- drivers
 - installing and configuring Emulex, 57

- dynamic disks
 - using, 68

E

- Emulex HBAs
 - installing and configuring, 57
 - installing and configuring drivers, 57

F

- failback
 - about, 106
- failover
 - configuring, 93
 - modifying failover policy, 106
 - policies, 99
- failover configurations
 - for Windows 2000 Advanced Server, 99
 - for Windows Server 2003, Enterprise Edition, 99

- failover policies, 99
 - failover pair, 102
 - failover ring, 104
 - for Windows 2000 Advanced Server cluster configurations, 99
 - for Windows Server 2003, Enterprise Edition, 99
 - multiway failover, 103
 - N+I failover, 100

- Fibre Channel
 - about, 15
 - switch fabric, 16

- file share resource type, 96

G

- group affinity
 - about, 102
 - configuring, 102

H

- HBA drivers
 - installing and configuring, 57
- high availability
 - about, 9
- host bus adapter
 - configuring the Fibre Channel HBA, 57

I

- IP address
 - assigning to cluster resources and components, 52
 - example configuration, 53

K

- keyboard
 - cabbling, 23

L

- LUNs
 - assigning to hosts, 69
 - configuring and managing, 68

M

- Microsoft Cluster Administrator
 - about, 83
 - running on a cluster node, 83
- MirrorView
 - about, 12, 14
- monitor
 - cabbling, 23
- mouse
 - cabbling, 23
- MSCS
 - installing and configuring, 74
 - uninstalling, 115
 - verifying operation, 80
- multiway failover, 103

N

- N+I failover
 - configuring group affinity, 100
- Navisphere Agent
 - about, 14, 65
- Navisphere Manager
 - about, 12-13, 64
 - hardware view, 12
 - storage view, 12
- network adapters
 - cabling the private network, 27-28
 - cabling the public network, 27
 - using dual-port for the private network, 56
- network failure
 - preventing, 87
- network interfaces, 88
- networking
 - configuring Windows, 52

O

- operating system
 - installing, 47, 50
 - rolling upgrade, 108
 - standard upgrade, 108
 - upgrading, 108

P

- period values
 - adjusting, 93
- power supplies
 - cabling, 23
- PowerPath
 - about, 13, 65
- private network
 - cabling, 26, 28
 - configuring IP addresses, 53
 - creating separate subnets, 54
 - hardware components, 28
 - hardware components and connections, 28
 - using dual-port network adapters, 56
- public network
 - cabling, 26
 - creating separate subnets, 54

Q

- quorum resource
 - definition, 10
- quorum resource
 - about, 10, 92
 - creating a LUN, 78
 - installing, 78
 - preventing failure, 78
 - running chkdsk, 116

R

RAID

- configuring the RAID level, 68

resource

- creating, 94
- deleting, 95

- resource dependencies, 90, 94

resource groups, 9

- definition, 9

- resource properties, 91

S

SAN

- about, 16
- configuring SAN backup in your cluster, 45

SAN-attached cluster

- about, 35
- configurations, 17

shared storage

- assigning drive letters, 69
- assigning LUNs to hosts, 69
- naming and formatting drives, 69

single initiator zoning

- about, 59

SnapView

- about, 12, 14

software

- for storage management, 13

storage groups

- about, 63

storage management software

- Access Control, 62
- Access Logix, 61
- MirrorView, 14
- Navisphere Agent, 14, 65
- Navisphere Manager, 13, 64
- PowerPath, 13, 65
- SnapView, 14

storage system

- configuring and managing LUNs, 68
- configuring drives on multiple shared storage systems, 72
- configuring the hard drives, 67
- using dynamic disks and volumes, 68

subnets

- creating, 54

- system requirements, 18

T

tape library

- connecting to a PowerEdge cluster, 44

threshold

- adjusting, 93

troubleshooting

- connecting to a cluster, 124
- failed resources, 81
- shared storage subsystem, 122

U

- upgrading
 - operating system, 108
- upgrading to a cluster solution
 - before you begin, 107
 - completing the upgrade, 107

V

- virtual servers, 9
 - definition, 9
- volumes
 - using, 68

W

- warranty, 21
- Windows 2000 Advanced Server
 - cluster configurations, 99
 - installing, 47
- Windows Server 2003,
Enterprise Edition
 - cluster configurations, 100,
102-104
 - installing, 47
- worldwide port name zoning, 58

Z

- zones
 - about, 16
 - implementing on a Fibre Channel
switched fabric, 57
 - in SAN configurations, 58
 - using worldwide port names, 58