

Dell|EMC CX3-series
Fibre Channel Storage Arrays With
Microsoft[®] Windows Server[®]
Failover Clusters

**Hardware 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.

© 2008 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 *Dell OpenManage* are trademarks of Dell Inc.; *Active Directory*, *Microsoft*, *Windows*, *Windows Server*, and *Windows NT* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.; *EMC*, *EMC ControlCenter*, *Navisphere*, and *PowerPath* are registered trademarks and *Access Logix*, *MirrorView*, *SAN Copy*, and *SnapView* are trademarks of EMC 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.

Contents

1	Introduction	7
	Cluster Solution	8
	Cluster Hardware Requirements	8
	Cluster Nodes.	9
	Cluster Storage	10
	Supported Cluster Configurations	12
	Direct-Attached Cluster	12
	SAN-Attached Cluster	13
	Other Documents You May Need	14
2	Cabling Your Cluster Hardware	17
	Cabling the Mouse, Keyboard, and Monitor	17
	Cabling the Power Supplies	17
	Cabling Your Cluster for Public and Private Networks	19
	Cabling the Public Network	20
	Cabling the Private Network	21
	NIC Teaming	21
	Cabling the Storage Systems	22
	Cabling Storage for Your Direct-Attached Cluster	22
	Cabling Storage for Your SAN-Attached Cluster	27

3	Preparing Your Systems for Clustering	39
	Cluster Configuration Overview	39
	Installation Overview	41
	Installing the Fibre Channel HBAs	42
	Installing the Fibre Channel HBA Drivers.	42
	Implementing Zoning on a Fibre Channel Switched Fabric	42
	Using Zoning in SAN Configurations Containing Multiple Hosts	43
	Using Worldwide Port Name Zoning	43
	Installing and Configuring the Shared Storage System	45
	Access Logix	45
	Access Control	47
	Storage Groups	47
	Navisphere Manager	49
	Navisphere Agent.	49
	EMC PowerPath.	49
	Enabling Access Logix and Creating Storage Groups Using Navisphere 6.x	50
	Configuring the Hard Drives on the Shared Storage System(s)	51
	Optional Storage Features	55
	Updating a Dell EMC Storage System for Clustering	56
	Installing and Configuring a Failover Cluster	56

A	Troubleshooting	57
B	Cluster Data Form	63
C	Zoning Configuration Form	65

Introduction

A Dell™ Failover Cluster combines specific hardware and software components to provide enhanced availability for applications and services that are run on the cluster. A Failover Cluster is designed to reduce the possibility of any single point of failure within the system that can cause the clustered applications or services to become unavailable. It is recommended that you use redundant components like server and storage power supplies, connections between the nodes and the storage array(s), and connections to client systems or other servers in a multi-tier enterprise application architecture in your cluster.

This document provides information to configure your Dell | EMC CX3-series fibre channel storage array with one or more Failover Clusters. It provides specific configuration tasks that enable you to deploy the shared storage for your cluster.

For more information on deploying your cluster with Windows Server 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com. For more information on deploying your cluster with Windows Server 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

For a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Failover Cluster, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

Cluster Solution

Your cluster implements a minimum of two node clustering to a maximum of either eight nodes (for Windows Server 2003) or sixteen nodes (for Windows Server 2008) clustering and provides the following features:

- 8-Gbps, 4-Gbps, and 2-Gbps 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

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 Small Computer System Interface (SCSI). Because Fibre Channel devices are hot-pluggable, you can add or remove devices from the nodes without taking the entire cluster offline.

Cluster Hardware Requirements

Your cluster requires the following hardware components:

- Servers (cluster nodes)
- Storage Array and storage management software

Cluster Nodes

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

Table 1-1. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	A minimum of two identical PowerEdge servers are required. The maximum number of nodes that is supported depends on the variant of the Windows Server operating system used in your cluster, and on the physical topology in which the storage system and nodes are interconnected.
RAM	The variant of the Windows Server operating system that is installed on your cluster nodes determines the minimum required amount of system RAM.
HBA ports	Two Fibre Channel HBAs per node, unless the server employs an integrated or supported dual-port Fibre Channel HBA. Where possible, place the HBAs on separate PCI buses to improve availability and performance.
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 disk controller. 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.



NOTE: For more information about supported systems, HBAs and operating system variants, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

Cluster Storage

Table 1-2 lists supported storage systems and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems.

Table 1-2. Cluster Storage Requirements

Hardware Components	Requirement
Supported storage systems	One to four supported Dell EMC storage systems. See Table 1-3 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 "Installing and Configuring the Shared Storage System" on page 45.

Table 1-3 lists hardware requirements for the storage processor enclosures (SPE), disk array enclosures (DAE), and standby power supplies (SPS).

Table 1-3. Dell|EMC Storage System Requirements

Processor Enclosure	Minimum Storage	Possible Storage Expansion	SPS
CX3-10c SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to three DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS
CX3-20c SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to seven DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS
CX3-20f SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to seven DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS
CX3-40c SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to 15 DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS

Table 1-3. Dell|EMC Storage System Requirements (continued)

Processor Enclosure	Minimum Storage	Possible Storage Expansion	SPS
CX3-40f SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to 15 DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS
CX3-80 SPE	One DAE3P-OS with at least five and up to 15 hard drives	Up to 31 DAE with a maximum of 15 hard drives each	Two per SPE and DAE3P-OS



NOTE: The DAE3P-OS is the first DAE enclosure that is connected to the CX3-series (including all of the storage systems listed above). Core software is preinstalled on the first five hard drives of the DAE3P-OS.

Each storage system in the cluster is centrally managed by one host system (also called a *management station*) running EMC Navisphere® Manager—a centralized storage management application used to configure Dell|EMC storage systems. Using a graphical user interface (GUI), you can select a specific view of your storage arrays, as shown in Table 1-4.

Table 1-4. 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 local area network (LAN) bandwidth.

For more information about Navisphere Manager, EMC Access Logix™, MirrorView, SnapView, and SAN Copy, see "Installing and Configuring the Shared Storage System" on page 45.

Supported Cluster Configurations

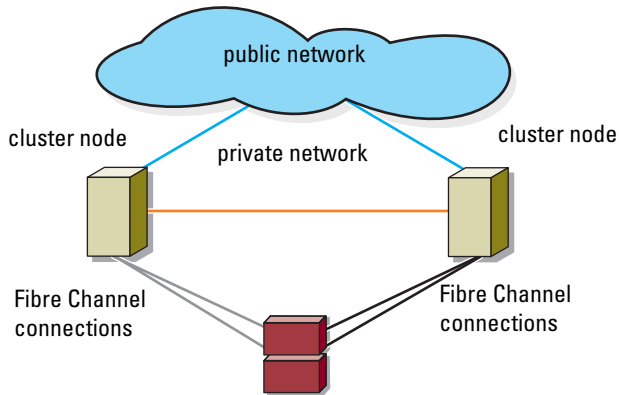
The following sections describe the supported cluster configurations.

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, multiple path detection, and 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 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




Other Documents You May Need

CAUTION: 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 *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.
- For more information on deploying your cluster with Windows Server 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide*.

- For more information on deploying your cluster with Windows Server 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide*.
 - 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.
 - Any other documentation that came with your server or storage system.
 - The EMC PowerPath documentation that came with your HBA kit(s) and Dell|EMC Storage Enclosure User's Guides.
 - 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 connections of each node to the switch box.

Cabling the Power Supplies


See the documentation for each component in your cluster solution and 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.

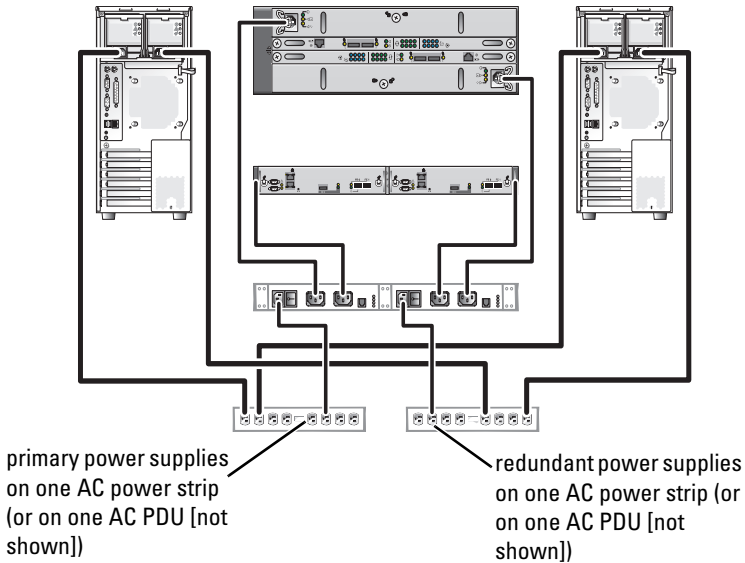
Figure 2-1 and Figure 2-2 illustrate recommended methods for power cabling for a cluster solution consisting of two PowerEdge systems and two storage systems. To ensure redundancy, the primary power supplies of all the components are grouped into one or two circuits and the redundant power supplies are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

Figure 2-1. Power Cabling Example With One Power Supply in the PowerEdge Systems



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

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



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

Cabling Your Cluster for Public and Private Networks

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


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.

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 <i>Mixed mode</i> for private network failover.
Private network	A dedicated connection for sharing cluster health and status information only.

Figure 2-3 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 2-3. 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: Throughout this document, *Gigabit Ethernet* is used to refer to either Gigabit Ethernet or 10 Gigabit Ethernet.

Using Dual-Port Network Adapters

You can configure your cluster to use the public network as a failover for private network communications. If you are using dual-port network adapters, do not configure both ports simultaneously to support both 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, only in a public network. NIC teaming is not supported in a private network.

Use the same brand of NICs in a team. Do not mix brands in NIC teaming.

Cabling the Storage Systems


This section provides information on cabling 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 host bus adapter (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 2-4 shows an example of a direct-attached, single cluster configuration with redundant HBA ports installed in each cluster node.

Figure 2-4. Direct-Attached Cluster Configuration



Cabling a Cluster to a Dell|EMC Storage System

Each cluster node attaches to the storage system using two Fibre optic cables with duplex local connector (LC) multimode connectors that attach to the HBA ports in the cluster nodes and the storage processor (SP) ports in the Dell|EMC storage system. These connectors consist of two individual Fibre 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 Dell|EMC 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 Fibre port 0 (first fibre channel port).
 - b Install a cable from cluster node 1 HBA port 1 to SP-B Fibre port 0 (first fibre channel port).
- 2 Connect cluster node 2 to the storage system:
 - a Install a cable from cluster node 2 HBA port 0 to SP-A Fibre port 1 (second fibre channel port).
 - b Install a cable from cluster node 2 HBA port 1 to SP-B Fibre port 1 (second fibre channel port).

Figure 2-5, Figure 2-6, and Figure 2-7 illustrate methods of cabling a two-node direct-attached cluster to a CX3-10c, CX3-20, and CX3-40f 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 2-5. Cabling the Cluster Nodes to a CX3-10c Storage System




Figure 2-6. Cabling the Cluster Nodes to a CX3-20 Storage System

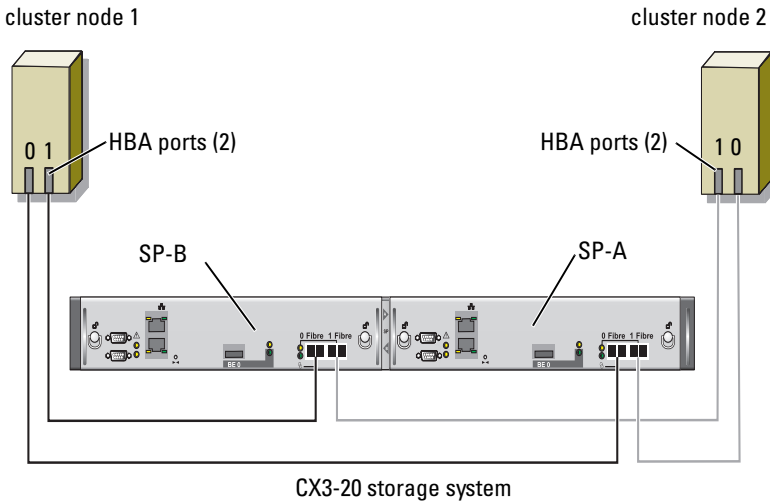
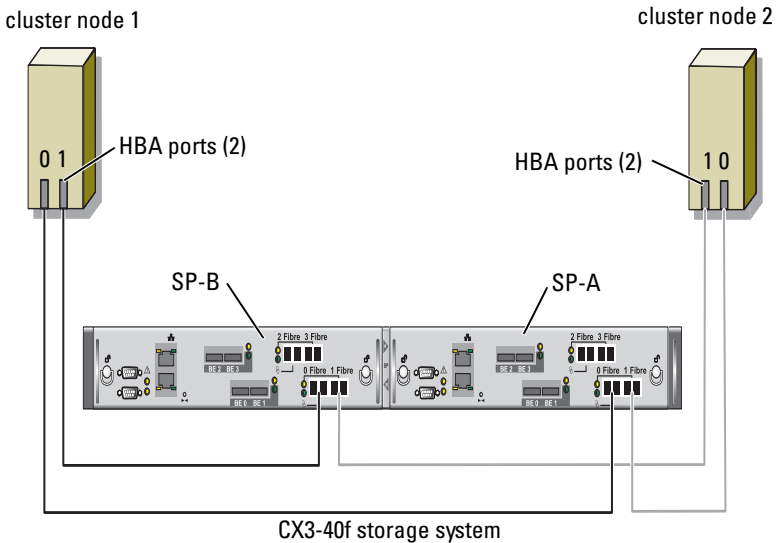


Figure 2-7. Cabling the Cluster Nodes to a CX3-40f Storage System



NOTE: If your cluster is attached to Dell|EMC CX3-10c, CX3-20/c, or CX3-40/c storage systems, you can configure two cluster nodes in a direct-attached configuration. With the CX3-40f or CX3-80 you can configure four cluster nodes and with CX3-20f you can configure six cluster nodes in a direct attached configuration.

Cabling a Four-Node Cluster to a CX3-20f, CX3-40f, and CX3-80 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 Fibre port 0 (first Fibre Channel port).
 - b** Install a cable from cluster node 1 HBA port 1 to SP-B Fibre port 0 (first Fibre Channel port).
- 2** Connect cluster node 2 to the storage system:
 - a** Install a cable from cluster node 2 HBA port 0 to SP-A Fibre port 1 (second Fibre Channel port).
 - b** Install a cable from cluster node 2 HBA port 1 to SP-B Fibre port 1 (second Fibre Channel port).

- 3 Connect cluster node 3 to the storage system:
 - a Install a cable from cluster node 3 HBA port 0 to SP-A Fibre port 2 (third Fibre Channel port).
 - b Install a cable from cluster node 3 HBA port 1 to SP-B Fibre port 2 (third Fibre Channel port).
- 4 Connect cluster node 4 to the storage system:
 - a Install a cable from cluster node 4 HBA port 0 to SP-A Fibre port 3 (fourth Fibre Channel port).
 - b Install a cable from cluster node 4 HBA port 1 to SP-B Fibre port 3 (fourth Fibre Channel port).

Cabling Two Clusters to a Dell|EMC Storage System

The four fibre channel ports per storage processor on Dell|EMC CX3-40f and CX3-80 storage systems allows you to connect two two-node clusters in a direct-attached configuration. Similarly, the six fibre-channel ports per processor on Dell|EMC CX3-20f storage system allows you to connect three two-node clusters in a direct-attached environment.



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

Cabling Two Two-Node Clusters to a CX3-20f, CX3-40f, or CX3-80 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 Fibre port 0 (first Fibre Channel port).
 - b Install a cable from cluster node 1 HBA port 1 to SP-B Fibre port 0 (first Fibre Channel port).
- 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 Fibre port 1 (second Fibre Channel port).
 - b Install a cable from cluster node 2 HBA port 1 to SP-B Fibre port 1 (second Fibre Channel port).

- 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 Fibre port 2 (third Fibre Channel port).
 - b Install a cable from cluster node 1 HBA port 1 to SP-B Fibre port 2 (third Fibre Channel port).
- 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 Fibre port 3 (fourth Fibre Channel port).
 - b Install a cable from cluster node 2 HBA port 1 to SP-B Fibre port 3 (fourth Fibre Channel port).

Cabling Storage for Your SAN-Attached Cluster

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

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

See "Implementing Zoning on a Fibre Channel Switched Fabric" on page 42 for more information on Fibre Channel switch fabrics.

Figure 2-8 shows an example of a two node SAN-attached cluster.

Figure 2-9 shows an example of an eight-node SAN-attached cluster.

Similar cabling concepts can be applied to clusters that contain a different number of nodes.



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 2-8. Two-Node SAN-Attached Cluster





Figure 2-9. Eight-Node SAN-Attached Cluster



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

The supported Dell|EMC storage systems are configured with one storage processor enclosure (SPE), at least one disk array enclosure (DAE) enclosure, and two standby power supplies (SPSs).

The cluster nodes attach to the storage system using a redundant switch fabric and Fibre 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 fibre 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 19 for more information on the duplex LC multimode fibre optic connector.

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



NOTE: Dell|EMC CX3-20f, CX3-40f, and CX3-80 storage systems have more than two fibre-channel ports per SP. You can also connect the additional ports to the redundant fabrics to achieve higher availability.

Table 2-3 provides information for cabling your storage system to the Fibre Channel switch.

Table 2-3. Storage System Cabling Description

Storage System	SP Ports	Fibre Optic Cables Required	Cabling Description
CX3-10c, CX3-20/c, CX3-40/c	Two ports per SP	4	Attach one cable from each storage processor port to the Fibre Channel switch.
CX3-40f, CX3-80	Four ports per SP	At least 4 and up to 8	
CX3-20f	Six ports per SP	At least 4 and up to 12	

Figure 2-10 and Figure 2-11 illustrate methods for cabling a SAN-attached cluster to the CX3-20 and CX3-40c storage systems, respectively.

Figure 2-12 illustrates a method for cabling a SAN-attached cluster to a CX3-80 storage system.

Cabling a SAN-Attached Cluster to a Dell|EMC CX3-10c, CX3-20/c, or CX3-40/c 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 Fibre port 0 (first Fibre Channel port).
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 1 (second Fibre Channel port).
 - c** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 1 (second Fibre Channel port).
 - d** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 0 (first Fibre Channel port).

Figure 2-10. Cabling a SAN-Attached Cluster to the Dell|EMC CX3-20 SPE

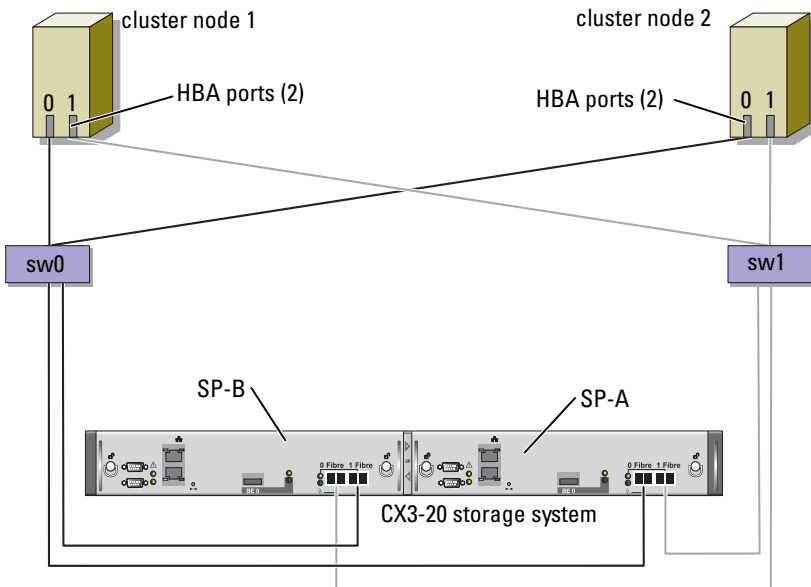



Figure 2-11. Cabling a SAN-Attached Cluster to the CX3-40c SPE




Cabling a SAN-Attached Cluster to the CX3-40 or CX3-80 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 Fibre port 0 (first Fibre Channel port).
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A Fibre port 2 (third Fibre Channel port).
 - c** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 1 (second Fibre Channel port).

- d Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 3 (fourth Fibre Channel port).
- e Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 1 (second Fibre Channel port).
- f Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 3 (fourth Fibre Channel port).
- g Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 0 (first Fibre Channel port).
- h Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 2 (third Fibre Channel port).

Figure 2-12. Cabling a SAN-Attached Cluster to the CX3-80



A CX3-20f storage system with six Fibre Channel ports per SP can be cabled in similar manner by connecting the remaining Fibre Channel ports to the switches if required.

Cabling 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.

For rules and guidelines for SAN-attached clusters, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.



NOTE: The following procedures use Figure 2-10, Figure 2-11, and Figure 2-12 as examples for cabling additional clusters.

Cabling Multiple SAN-Attached Clusters to the CX3-10c, CX3-20, or CX3-40c 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 Fibre port 0 (first Fibre Channel port).
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 1 (second Fibre Channel port).
 - c** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 1 (second Fibre Channel port).
 - d** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 0 (first Fibre Channel port).

Cabling Multiple SAN-Attached Clusters to the CX3-40f or CX3-80 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 Fibre port 0 (first Fibre Channel port).
 - b** Connect a cable from Fibre Channel switch 0 (sw0) to SP-A Fibre port 2 (third Fibre Channel port).
 - c** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 1 (second Fibre Channel port).
 - d** Connect a cable from Fibre Channel switch 0 (sw0) to SP-B Fibre port 3 (fourth Fibre Channel port).
 - e** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 1 (second Fibre Channel port).
 - f** Connect a cable from Fibre Channel switch 1 (sw1) to SP-A Fibre port 3 (fourth Fibre Channel port).
 - g** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 0 (first Fibre Channel port).
 - h** Connect a cable from Fibre Channel switch 1 (sw1) to SP-B Fibre port 2 (third Fibre Channel port).



NOTE: Dell|EMC CX3-20f storage system can be cabled in a similar manner as the CX3-40f or CX3-80 storage system. The remaining fibre channel ports (4Fibre and 5Fibre) in the CX3-20f storage system can also be connected depending upon the required level of redundancy.

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 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. Failover Clusters can support configurations with multiple storage units attached to clustered nodes. In this scenario, the Microsoft Cluster Service (MSCS) software can fail over disk drives in any cluster-attached shared storage array between the cluster nodes.



NOTE: Throughout this document, MSCS is used to refer to either the Microsoft Windows Server 2003 Cluster Service or the Microsoft Windows Server 2008 Failover Cluster Service.


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 "If you have Access Control enabled in Navisphere Manager, you must create storage groups and assign LUNs to the proper host systems." on page 53 for more information.

Figure 2-13 provides an example of cabling the cluster nodes to four Dell|EMC storage systems. See "Implementing Zoning on a Fibre Channel Switched Fabric" on page 42 for more information.

Figure 2-13. 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 may contain an integrated Fibre Channel bridge or Storage Network Controller (SNC) that connects directly to your Dell|EMC Fibre Channel switch.

Figure 2-14 shows a supported Failover 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 2-14. 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 2-15 provides an example of cabling the cluster nodes to your storage systems and SAN backup with a tape library.

Figure 2-15. Cluster Configuration Using SAN-Based Backup




Preparing Your Systems for Clustering


 **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.

Cluster Configuration Overview

- 1 Ensure that your site can handle the cluster's power requirements. Contact your sales representative for information about your region's power requirements.
- 2 Install the systems, the shared storage array(s), and the interconnect switches (for example, in an equipment rack), and ensure that all the components are turned on.

 **NOTE:** For more information on step 3 to step 7 and step 10 to step 13, see the "Preparing your systems for clustering" section of *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* or *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

- 3 Deploy the operating system (including any relevant service packs and hotfixes), network adapter drivers, and storage adapter drivers (including Multipath I/O (MPIO) drivers) on each cluster node. Depending on the deployment method that is used, it may be necessary to provide a network connection to successfully complete this step.

 **NOTE:** To help in planning and deployment of your cluster, record the relevant cluster configuration information in the Cluster Data Form located at "Cluster Data Form" on page 63 and the Zoning configuration information in the Zoning Configuration form located and "Zoning Configuration Form" on page 65.

- 4 Establish the physical network topology and the TCP/IP settings for network adapters on each cluster node to provide access to the cluster public and private networks.

- 5 Configure each cluster node as a member in the same Windows Active Directory Domain.



NOTE: You can configure the cluster nodes as Domain Controllers. For more information, see the "Selecting a Domain Model" section of *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* or *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

- 6 Establish the physical storage topology and any required storage network settings to provide connectivity between the storage array and the systems that you are configuring as cluster nodes. Configure the storage system(s) as described in your storage system documentation.
- 7 Use storage array management tools to create at least one logical unit number (LUN). The LUN is used as a cluster Quorum disk for Windows Server 2003 Failover cluster and as a Witness disk for Windows Server 2008 Failover cluster. Ensure that this LUN is presented to the systems that you are configuring as cluster nodes.



NOTE: For security reasons, it is recommended that you configure the LUN on a single node as mentioned in step 8 when you are setting up the cluster. Later, you can configure the LUN as mentioned in step 9 so that other nodes in the cluster can access it.

- 8 Select one of the systems and form a new failover cluster by configuring the cluster name, cluster management IP, and quorum resource. For more information, see "Preparing Your Systems for Clustering" on page 39.



NOTE: For Failover Clusters configured with Windows Server 2008, run the **Cluster Validation Wizard** to ensure that your system is ready to form the cluster.

- 9 Join the remaining node(s) to the failover cluster. For more information, see "Preparing Your Systems for Clustering" on page 39.
- 10 Configure roles for cluster networks. Take any network interfaces that are used for iSCSI storage (or for other purposes outside of the cluster) out of the control of the cluster.
- 11 Test the failover capabilities of your new cluster.



NOTE: For Failover Clusters configured with Windows Server 2008, you can also use the **Cluster Validation Wizard**.

- 12 Configure highly-available applications and services on your Failover Cluster. Depending on your configuration, this may also require providing additional LUNs to the cluster or creating new cluster resource groups. Test the failover capabilities of the new resources.
- 13 Configure client systems to access the highly-available applications and services that are hosted on your failover cluster.

Installation Overview

Each node in your Dell Failover Cluster must be installed with the same release, edition, service pack, and processor architecture of the Windows Server operating system. For example, all nodes in your cluster may be configured with Windows Server 2003 R2, Enterprise x64 Edition. If the operating system varies among nodes, it is not possible to configure a Failover Cluster successfully. It is recommended to establish server roles prior to configuring a Failover Cluster, depending on the operating system configured on your cluster.

For a list of Dell PowerEdge Servers, Fibre Channel HBAs and switches, and recommended list of operating system variants, specific driver and firmware revisions, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

For a general overview of cluster configuration tasks and more detailed information about deploying your cluster with Windows Server 2003 operating system, see the Dell Failover Clusters with *Microsoft Windows Server 2003 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

For more information on deploying your cluster with Windows Server 2008 operating systems, see the Dell Failover Clusters with *Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

The following sub-sections describe steps that must be taken to enable communication between the cluster nodes and your shared Dell | EMC CX3-series Fibre Channel storage array, and to present disks from the storage array to the cluster.

Installing the Fibre Channel HBAs

For dual-HBA configurations, it is recommended that you install the Fibre Channel HBAs on separate peripheral component interconnect (PCI) buses. Placing the adapters on separate buses improves availability and performance.

For more information about your system's PCI bus configuration and supported HBAs, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

Installing the Fibre Channel HBA Drivers

For more information, see the EMC documentation that is included with your HBA kit.

For more information about installing and configuring Emulex HBAs and EMC-approved drivers, see the Emulex support website located at www.emulex.com or the Dell Support website at support.dell.com.

For more information about installing and configuring QLogic HBAs and EMC-approved drivers, see the QLogic support website at www.qlogic.com or the Dell Support website at support.dell.com.

For more information about supported HBA controllers and drivers, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

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, storage processor (SP) ports, and Fibre Channel to SCSI bridges or storage network controllers (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 3-1 provides a list of WWN identifiers that you can find in the Dell|EMC cluster environment.

Table 3-1. 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
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	Dell 2362M HBA port

Table 3-1. Port Worldwide Names in a SAN Environment (continued)

Identifier	Description
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
xx:xx:xx:C0:97:xx:xx:xx	PowerVault ML6000 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 logical unit numbers (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 CX3-series 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/SNC or fibre-channel interface 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 "Cluster Hardware Requirements" on page 8 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 50 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 3-2 provides a list of cluster and host system configurations and their Access Logix requirement.

Table 3-2. 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.

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 3-3 describes the properties in the storage group.

Table 3-3. 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	<p>Lists the host systems connected to the storage group.</p> <p>Each host entry contains the following fields:</p> <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 <p>NOTE: In a clustered environment, all nodes of a cluster must be connected to the same storage group.</p>
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 graphical user interface (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 user interface (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

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**. It is recommended to turn off 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.

It is recommended that you create 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

For more information on deploying your cluster with Windows Server 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

For more information on deploying your cluster with Windows Server 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

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.



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. Format the LUNs as NTFS 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.

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.



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 Server 2003 Disk Management.
- 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 17 in "Cabling Your Cluster Hardware" on page 17 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 54.

Formatting and Assigning Drive Letters and Volume Labels to the Disks

- 1 Turn off 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 Turn off 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 Turn off 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.

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

Optional Storage Features

Your Dell|EMC CX3-series storage array may be configured to provide optional features that can be used in conjunction with your cluster. These features include MirrorView, SnapView, and SANCopy.

MirrorView

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

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.

The source LUN and each snapshot or clone must be accessed from a different host or a different cluster.

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

SAN Copy

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.

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.

Upgrade the core software version that is running on the storage system or enable Access Logix. For specific version requirements, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability website at www.dell.com/ha.

Installing and Configuring a Failover Cluster

After you have established the private and public networks and have assigned the shared disks from the storage array to the cluster nodes, you can configure the operating system services on your Dell Failover Cluster. The procedure to configure the Failover Cluster depends on the version of the Windows Server operating system that is running on the system.

For more information on deploying your cluster with Windows Server 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

For more information on deploying your cluster with Windows Server 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

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 19 for more information.
	The length of the interface cables exceeds the maximum allowable length.	Ensure that the fibre 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. One or more nodes may have the Internet Connection Firewall enabled, blocking Remote Procedure Call (RPC) communications between the nodes. 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. 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. 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.

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: Microsoft Cluster Service successfully formed a cluster on this node. or Microsoft Cluster Service successfully joined the cluster. 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.
	A cluster has not been formed on the system.	
	The system has just been booted and services are still starting.	
	The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes.	Configure the Internet Connection Firewall to allow communications that are required by 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
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. For more information about assigning the network IPs, see "Assigning Static IP Addresses to Cluster Resources and Components" of <i>Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide</i> or <i>Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide</i> .
	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 Server 2003 cluster generates error messages.	Some resources in Windows Server 2003 are not supported in Windows NT 4.0.	It is strongly recommended that you use Microsoft 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>One or more nodes may have the Internet Connection Firewall enabled, blocking RPC communications between the 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>Configure the Internet Connection Firewall to allow communications that are required by the 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	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.	No action required.

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 and the Internet Firewall enabled.	The Windows Internet Connection Firewall is enabled, which may conflict with Cluster Services.	Perform the following steps: <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>
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 RPC communications between the nodes.	Configure the Internet Connection Firewall to allow communications that are required by the 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 Information	Cluster Solution
Cluster name and IP address	
Server type	
Installer	
Date installed	
Applications	
Location	
Notes	

Table B-2. Cluster Node Information

Node Name	Service Tag Number	Public IP Address	Private IP Address

Additional Networks

Table B-3. Storage Array Information

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, 47
- Access Logix
 - about, 45

C

- cable configurations
 - cluster interconnect, 21
 - for client networks, 20
 - for mouse, keyboard, and monitor, 17
 - for power supplies, 17
- cluster
 - optional configurations, 12
- cluster configurations
 - connecting to multiple shared storage systems, 36
 - connecting to one shared storage system, 12
 - direct-attached, 12, 22
 - SAN-attached, 13
- cluster storage
 - requirements, 10
- clustering
 - overview, 7

D

- Dell | EMC CX3-40
 - Cabling a two-node cluster, 23
- Dell | EMC CX3-80
 - cabling to two clusters, 26
- Dell | EMC CX3-20
 - Cabling a two-node cluster, 23
- Dell | EMC CX3-80
 - cabling the cluster nodes, 23
 - cabling the cluster nodes in a SAN-attached environment, 30
 - cabling to one SAN-attached cluster, 30
 - configuring, 45
 - installing, 45
 - updating for cluster use, 56
 - zoning in a switched environment, 35
- direct-attached cluster
 - about, 22
 - cabling, 23
- drivers
 - installing and configuring Emulex, 42
- dynamic disks
 - using, 52

E

- Emulex HBAs
 - installing and configuring, 42
 - installing and configuring drivers, 42

H

- HBA drivers
 - installing and configuring, 42
- host bus adapter
 - configuring the Fibre Channel HBA, 42

K

- keyboard
 - cabling, 17

L

- LUNs
 - assigning to hosts, 53
 - configuring and managing, 52

M

- MirrorView
 - about, 12
- monitor
 - cabling, 17

- mouse
 - cabling, 17

- MSCS
 - installing and configuring, 56

N

- Navisphere Agent
 - about, 49
- Navisphere Manager
 - about, 11, 49
 - hardware view, 11
 - storage view, 11
- network adapters
 - cabling the private network, 20-21
 - cabling the public network, 20

O

- operating system
 - Windows Server 2003, Enterprise Edition
 - installing, 41

P

- power supplies
 - cabling, 17
- PowerPath
 - about, 49

- private network
 - cabling, 19, 21
 - hardware components, 21
 - hardware components and connections, 21
- public network
 - cabling, 19

R

- RAID
 - configuring the RAID level, 52

S

- SAN
 - configuring SAN backup in your cluster, 38
- SAN-attached cluster
 - about, 27
 - configurations, 12
- shared storage
 - assigning LUNs to hosts, 53
 - naming and formatting drives, 53
- single initiator zoning
 - about, 44
- SnapView
 - about, 12
- storage groups
 - about, 47

- storage management software
 - Access Control, 47
 - Access Logix, 45
 - Navisphere Agent, 49
 - Navisphere Manager, 49
 - PowerPath, 49
- storage system
 - configuring and managing LUNs, 52
 - configuring drives on multiple shared storage systems, 53
 - configuring the hard drives, 51
 - using dynamic disks and volumes, 52

T

- tape library
 - connecting to a PowerEdge cluster, 37
- troubleshooting
 - connecting to a cluster, 59
 - shared storage subsystem, 57

V

- volumes
 - using, 52

W

warranty, 14

worldwide port name zoning, 43

Z

zones

- implementing on a Fibre Channel
switched fabric, 42

- in SAN configurations, 43

- using worldwide port names, 43