

Dell/EMC CX4-series  
iSCSI Storage Arrays With  
Microsoft® Windows Server®  
Failover Clusters

# Hardware Installation and Troubleshooting Guide



# Notes, Cautions, and Warnings



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



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



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

---

**Information in this document is subject to change without notice.**

**© 2009–2010 Dell Inc. All rights reserved.**

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc.; *Active Directory*, *Microsoft*, *Windows*, *Windows Server*, *Windows XP* and *Windows NT* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries; *EMC*, *Navisphere*, and *PowerPath* are registered trademarks and *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
	<b>Cluster Solution</b> . . . . .	<b>8</b>
	<b>Cluster Hardware Requirements</b> . . . . .	<b>8</b>
	Cluster Nodes. . . . .	8
	Cluster Storage . . . . .	9
	NICs Dedicated to iSCSI . . . . .	11
	Network Switches Dedicated to iSCSI. . . . .	11
	<b>Supported Cluster Configurations</b> . . . . .	<b>12</b>
	Direct-Attached Cluster . . . . .	12
	iSCSI Switch-Attached Cluster. . . . .	13
	<b>Other Documents You May Need</b> . . . . .	<b>13</b>
2	Cabling Your Cluster Hardware . . . . .	15
	<b>Cabling the Mouse, Keyboard, and Monitor</b> . . . . .	<b>15</b>
	<b>Cabling the Power Supplies</b> . . . . .	<b>15</b>
	<b>Cabling Your Cluster for Public     and Private Networks</b> . . . . .	<b>17</b>
	Cabling the Public Network . . . . .	19
	Cabling the Private Network . . . . .	19
	NIC Teaming . . . . .	20
	<b>Cabling the Storage Systems</b> . . . . .	<b>20</b>
	Cabling Storage for Your Direct-Attached Cluster . . . . .	20

	Cabling Storage for Your iSCSI Switch-Attached Cluster . . . . .	26
<b>3</b>	<b>Preparing Your Systems for Clustering . . . . .</b>	<b>37</b>
	<b>Cluster Configuration Overview . . . . .</b>	<b>37</b>
	<b>Installation Overview . . . . .</b>	<b>39</b>
	<b>Installing the iSCSI NICs. . . . .</b>	<b>40</b>
	Installing the Microsoft iSCSI Software Initiator . . . . .	40
	Modifying the TCP Registry Settings . . . . .	41
	<b>Installing and Configuring the Shared Storage System . . . . .</b>	<b>41</b>
	Access Control . . . . .	42
	Storage Groups . . . . .	42
	Navisphere Manager . . . . .	43
	Navisphere Agent. . . . .	44
	EMC PowerPath. . . . .	45
	Enabling Access Control and Creating Storage Groups Using Navisphere . . . . .	45
	Configuring the Hard Drives on the Shared Storage System(s) . . . . .	46
	Optional Storage Features . . . . .	47
	<b>Updating a Dell/EMC Storage System for Clustering. . . . .</b>	<b>48</b>
	<b>Installing and Configuring a Failover Cluster . . . . .</b>	<b>49</b>

A	Troubleshooting . . . . .	51
B	iSCSI Configuration Worksheet . . . . .	59
C	Cluster Data Form . . . . .	61



# 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 system and storage power supplies, connections between the nodes and the storage array(s), and connections to client systems or other systems in a multi-tier enterprise application architecture in your cluster.

This document provides information to configure your Dell/EMC CX4-series iSCSI storage array(s) 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 Microsoft® Windows Server® 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* available at [support.dell.com/manuals](http://support.dell.com/manuals).



**NOTE:** Throughout this document, Windows Server 2008 refers to both Windows Server 2008 and Windows Server 2008 R2.

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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

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 Clustering website at [dell.com/ha](http://dell.com/ha).

# Cluster Solution

For Windows Server 2003 operating system, your cluster implements a minimum of two to a maximum of eight node clustering. For Windows Server 2008 operating system, your cluster implements a minimum of two to a maximum of sixteen node clustering. The cluster solution provides the following features:

- Gigabit and 10 Gigabit Ethernet iSCSI technologies
- 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.

## Cluster Hardware Requirements

Your cluster requires the following hardware components:

- Cluster nodes
- Cluster storage

### 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 Dell PowerEdge™ systems are required. The maximum number of nodes that are supported depend 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 RAM required.
iSCSI Initiator	Install the iSCSI port driver, Initiator Service, and Software Initiator on each node.



**Table 1-1. Cluster Node Requirements (continued)**

Component	Minimum Requirement
Network Interface	Two iSCSI NICs or two iSCSI NIC ports per node: Configure the NICs on separate PCI buses to improve availability and iSCSI access.  TCP/IP Offload Engine (TOE) NICs are also supported for iSCSI traffic.
NICs	At least two NICs: one NIC for the public network and another NIC for the private network.  <b>NOTE:</b> 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).  <b>NOTE:</b> 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 and operating system variants, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability Clustering website at [dell.com/ha](http://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. See "Installing and Configuring the Shared Storage System" on page 41.

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**

<b>Processor Enclosure</b>	<b>Minimum Storage</b>	<b>Possible Storage Expansion</b>	<b>SPS</b>
CX4-120	One DAE-OS with at least five and up to 15 hard drives	Up to seven DAE's	Two for SPE and DAE-OS
CX4-240	One DAE-OS with at least five and up to 15 hard drives	Up to 15 DAE's	Two for SPE and DAE-OS
CX4-480	One DAE-OS with at least five and up to 15 hard drives	Up to 31 DAE's	Two for SPE and DAE-OS
CX4-960	One DAE-OS with at least five and up to 15 hard drives	Up to 63 DAE's	Two for SPE and DAE-OS

 **NOTE:** The DAE-OS is the first DAE enclosure that is connected to the CX4-series (including all the storage systems listed above). Core software is preinstalled on the first five hard drives of the DAE-OS.

 **NOTE:** The Flare version for 10 Gigabit iSCSI is 04.29 or later.

Each storage system in a cluster is centrally managed by one host system (also known as *management station*) running EMC® Navisphere® Manager. Navisphere Manager is a centralized storage management application used to configure Dell/EMC storage systems. You can select a specific view of your storage arrays using a web user interface as shown in Table 1-4.

**Table 1-4. Navisphere Manager Storage Views**

<b>View</b>	<b>Description</b>
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).

**Table 1-4. Navisphere Manager Storage Views (continued)**

<b>View</b>	<b>Description</b>
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 Navisphere Taskbar to download 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, MirrorView, SnapView, and SAN Copy, see "Installing and Configuring the Shared Storage System" on page 41.

## **NICs Dedicated to iSCSI**

The NIC controlled by iSCSI Software Initiator acts as an I/O adapter to connect the system's expansion bus and the storage components. Failover Cluster solutions that are configured with the CX4-series storage array require two iSCSI NICs or NIC ports in each PowerEdge system to provide redundant paths and load balance the I/O data transfer to the storage system.

## **Network Switches Dedicated to iSCSI**

The Gigabit or 10 Gigabit switch for iSCSI access functions as a regular network switch that provides extension and dedicated interconnection between the node and the storage system(s).



**NOTE:** It is recommended that you use dedicated switches for iSCSI traffic. To share the switches, use a separate network segment or virtual LAN (VLAN) for iSCSI traffic.

# Supported Cluster Configurations


The following sections describe the supported cluster configurations.

## Direct-Attached Cluster

In a direct-attached cluster, all the nodes of the cluster are directly attached to a single storage system. In this configuration, the RAID controllers (or storage processors) on the storage system are directly connected by cables to the Gigabit or 10 Gigabit Ethernet network interface cards (NICs) 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<sup>®</sup> provides failover capabilities, multiple path detection, and dynamic load balancing between multiple ports on the same storage processor. However, the 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.

## iSCSI Switch-Attached Cluster

In an iSCSI switch-attached cluster, all nodes are attached to a single storage system or to multiple storage systems through redundant LANs. iSCSI switch-attached clusters are superior to direct-attached clusters in configuration flexibility, expandability, and performance.

Figure 1-2 shows an iSCSI Switch-Attached Cluster.

**Figure 1-2. iSCSI Switch-Attached Cluster**



## Other Documents You May Need

**!** **WARNING:** The safety information that shipped with your system 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 available at [support.dell.com/manuals](http://support.dell.com/manuals).

- 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 *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* provides more information on deploying your cluster with the Windows Server 2003 operating systems.
- The *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* provides more information on deploying your cluster with the Windows Server 2008 operating systems.
- 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 system or storage system.
- The EMC PowerPath documentation 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 available at [support.dell.com/manuals](http://support.dell.com/manuals).

## 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 available at [support.dell.com/manuals](http://support.dell.com/manuals).

**Table 2-1. Network Connections**

<b>Network Connection</b>	<b>Description</b>
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.
iSCSI network	Two direct or switched iSCSI connections from the cluster node to the storage system.

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 two possible private network configurations.

**Table 2-2. Private Network Hardware Components and Connections**

Method	Hardware Components	Connection
Network switch	Gigabit or 10 Gigabit Ethernet network adapters and switches	Depending on the hardware, connect the CAT5e or CAT6 cables, the multi-mode optical cables with Local Connector (LC) connectors, or the twin-ax cables from the network adapters in the nodes to a switch.
Point-to-Point (two-node clusters only)	Copper Gigabit or 10 Gigabit Ethernet network adapters with RJ-45 connectors	Connect a <i>standard</i> CAT5e or CAT6 Ethernet cable between the network adapters in both nodes.
	Copper 10 Gigabit Ethernet network adapters with SFP+ connectors	Connect a twin-ax cable between the network adapters in both nodes.
	Optical Gigabit or 10 Gigabit Ethernet network adapters with LC connectors	Connect a multi-mode optical cable between the network adapters in both nodes.



**NOTE:** For more information on the supported cable types, see your system or NIC documentation.

## 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 an iSCSI switch-attached configuration.



**NOTE:** The illustrations in this section show examples for connecting the cluster nodes to the Gigabit iSCSI ports of the storage system. Use optical cables and the same cabling method to connect to the 10 Gigabit iSCSI ports of the storage system.




**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 Storage for Your Direct-Attached Cluster

A direct-attached cluster configuration consists of redundant iSCSI NIC ports cabled directly to a Dell/EMC storage system.

Figure 2-4 shows an example of a direct-attached, single cluster configuration with redundant iSCSI 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 Gigabit iSCSI ports of the storage system using CAT5e or CAT6 LAN cables with RJ45 connectors, or to the 10 Gigabit iSCSI ports of the storage system using multi-mode optical cables with LC connectors.

### Cabling a Two-Node Cluster to a Dell/EMC Storage System


**NOTE:** In direct-attached configurations, the 10 Gigabit iSCSI ports on the storage array can only be connected to the 10 Gigabit NICs on the servers.

- 1 Connect cluster node 1 to the storage system:
  - a Connect a LAN cable from NIC-0 of cluster node 1 to SP-A iSCSI port 0.
  - b Connect a LAN cable from NIC-1 of cluster node 1 to SP-B iSCSI port 0.
- 2 Connect cluster node 2 to the storage system:
  - a Connect a LAN cable from NIC-1 of cluster node 2 to SP-A iSCSI port 1.
  - b Connect a LAN cable from NIC-0 of cluster node 2 to SP-B iSCSI port 1.


**Figure 2-5. Cabling a Two-Node Cluster to a CX4-120 or CX4-240 Storage System**



**Figure 2-6. Cabling a Two-Node Cluster to a CX4-480 Storage System**



**Figure 2-7. Cabling a Two-Node Cluster to a CX4-960 Storage System**



### ***Cabling a Multi-Node Cluster to a Dell/EMC Storage System***

You can configure a cluster with more than two nodes in a direct-attached configuration using a Dell/EMC storage system, depending on the availability of iSCSI ports. With 1 Gigabit iSCSI, the CX4-120 can support up to 4-node cluster, CX4-240 and CX4-480 storage systems can support up to 6-node cluster, and the CX4-960 storage system can support up to 8-node cluster. With 10 Gigabit iSCSI, the CX4-120 and CX4-240 storage systems can support up to 2-node cluster and the CX-480 and CX-960 storage systems can support up to 4-node cluster.

The following example explains how to cable a four-node cluster:



**NOTE:** The Dell/EMC storage system requires at least 4 iSCSI ports available on each storage processor.



**NOTE:** The following steps can be modified to add more clusters or standalone systems.

- 1** Connect cluster node 1 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 1 to SP-A iSCSI port 0.
  - b** Connect a LAN cable from NIC-1 of cluster node 1 to SP-B iSCSI port 0.
- 2** Connect cluster node 2 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 2 to SP-A iSCSI port 1.
  - b** Connect a LAN cable from NIC-1 of cluster node 2 to SP-B iSCSI port 1.
- 3** Connect cluster node 3 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 3 to SP-A iSCSI port 2.
  - b** Connect a LAN cable from NIC-1 of cluster node 3 to SP-B iSCSI port 2.
- 4** Connect cluster node 4 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 4 to SP-A iSCSI port 3.



- b** Connect a LAN cable from NIC-1 of cluster node 4 to SP-B iSCSI port 3.

### **Cabling Multiple Clusters to a Dell/EMC Storage System**

The high number of available front-end iSCSI ports on the CX4-series storage system also allows you to configure multiple clusters or a mix of cluster(s) and non-clustered server(s) in a direct-attached configuration.

For example, four Gigabit iSCSI ports per storage processor on Dell/EMC CX4-120 storage or four 10 Gigabit iSCSI ports per storage processor on Dell/EMC CX4-480 and CX4-960 storage systems allow you to connect 2 two-node clusters, or 1 two-node clusters and two non-clustered systems in a direct-attached configuration. The six Gigabit iSCSI ports per storage processor on Dell/EMC CX4-240 and CX4-480 storage systems allows you to connect 3 two-node clusters, or 2 two-node clusters and two non-clustered systems in a direct-attached configuration. Similarly, the eight Gigabit iSCSI ports per storage processor on Dell/EMC CX4-960 storage system allows you to connect 4 two-node clusters, or 2 two-node clusters and four non-clustered systems in a direct-attached environment.



**NOTE:** Enable EMC® Access Control if the CX4-series storage system is connected to more than one cluster in a direct-attached configuration.

### ***Cabling Two Two-Node Clusters to a Dell/EMC Storage System***

The following steps are an example of how to cable 2 two-node clusters. The Dell/EMC storage system needs to have at least four iSCSI ports available on each storage processor.

- 1** In the first cluster, connect cluster node 1 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 1 to SP-A iSCSI port 0.
  - b** Connect a LAN cable from NIC-1 of cluster node 1 to SP-B iSCSI port 0.
- 2** In the first cluster, connect cluster node 2 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 2 to SP-A iSCSI port 1.
  - b** Connect a LAN cable from NIC-1 of cluster node 2 to SP-B iSCSI port 1.

- 3** In the second cluster, connect cluster node 1 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 1 to SP-A iSCSI port 2.
  - b** Connect a LAN cable from NIC-1 of cluster node 1 to SP-B iSCSI port 2.
- 4** In the second cluster, connect cluster node 2 to the storage system:
  - a** Connect a LAN cable from NIC-0 of cluster node 2 to SP-A iSCSI port 3.
  - b** Connect a LAN cable from NIC-1 of cluster node 2 to SP-B iSCSI port 3.

### **Cabling Storage for Your iSCSI Switch-Attached Cluster**


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

An iSCSI switch-attached cluster configuration provides more flexibility, expandability, and performance than a direct-attached configuration.


Figure 2-8 shows an example of a two-node, iSCSI switch-attached cluster. Each cluster node is configured with two iSCSI single channel network interface cards (NICs) or one iSCSI dual channel NIC that is connected to two Gigabit or 10 Gigabit Ethernet switches. The storage system includes two SPs that provide redundant connections to the network switches. This configuration provides data redundancy and eliminates a single point of failure on the host or storage system.

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


Figure 2-9 shows an example of an eight-node iSCSI switch-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 iSCSI Switch-Attached Cluster.**



**Figure 2-9. Eight-Node iSCSI Switch-Attached Cluster**



**Table 2-3. Storage System Cabling Description**

<b>Storage System</b>	<b>iSCSI ports per SP</b>	<b>Cables Required</b>	<b>Cabling Description</b>
CX4-120	Two to four Gigabit iSCSI ports	Four to eight CAT5e or CAT6 cables	Attach one cable from each storage processor port to the Gigabit or 10 Gigabit Ethernet switch.
	Up to two 10 Gigabit iSCSI ports	Up to four optical cables	
CX4-240	Two to six Gigabit iSCSI ports	Four to twelve CAT5e or CAT6 cables	
	Up to two 10 Gigabit iSCSI ports	Up to four optical cables	
CX4-480	Two to six Gigabit iSCSI ports	Four to twelve CAT5e or CAT6 cables	
	Up to four 10 Gigabit iSCSI ports	Up to eight optical cables	
CX4-960	Two to eight Gigabit iSCSI ports	Four to sixteen CAT5e or CAT6 cables	
	Up to four 10 Gigabit iSCSI ports	Up to eight optical cables	



**NOTE:** In switch-attached configurations, the Gigabit NICs can access the 10 Gigabit iSCSI ports on the storage system if:

- The switch supports both Gigabit and 10 Gigabit Ethernet
- The servers are connected to a Gigabit Ethernet switch, cascaded to the 10 Gigabit Ethernet that is attached to the storage system



**NOTE:** Adding more cables from the storage system to the switches can increase the I/O bandwidth and high availability of data.


Figure 2-10 illustrates the method for cabling an iSCSI switch-attached cluster to the CX4-120 and CX4-240 storage systems.

Figure 2-11 illustrates the method for cabling an iSCSI switch-attached cluster to a CX4-480 storage system.


Figure 2-12 illustrates the method for cabling an iSCSI switch-attached cluster to a CX4-960 storage system.

***Cabling an iSCSI Switch-Attached Cluster to a Dell/EMC CX4-120, CX4-240, CX4-480, or CX4-960 Storage System***


- 1 Connect cluster node 1 to the iSCSI network:
  - a Connect a LAN cable from NIC-0 to switch 0 (sw0).
  - b Connect a LAN cable from NIC-1 to switch 1 (sw1).
- 2 Repeat step 1 for each additional cluster node.
- 3 Connect the storage system to the iSCSI network:
  - a Connect a LAN cable from switch 0 (sw0) to SP-A iSCSI port 0.
  - b Connect a LAN cable from switch 0 (sw0) to SP-B iSCSI port 0.
  - c Connect a LAN cable from switch 1 (sw1) to SP-A iSCSI port 1.
  - d Connect a LAN cable from switch 1 (sw1) to SP-B iSCSI port 1.

 **NOTE:** Additional cables can be connected from the network switches to the storage system if there are available iSCSI ports on the storage processors.


**Figure 2-10. Cabling a iSCSI Switch-Attached Cluster to the Dell/EMC CX4-120 or CX4-240**



**Figure 2-11. Cabling an iSCSI-Attached Cluster to the Dell/EMC CX4-480**



**Figure 2-12. Cabling a iSCSI Switch-Attached Cluster to the Dell/EMC CX4-960**



### **Cabling Multiple iSCSI Switch-Attached Clusters to a Dell/EMC Storage System**

To cable multiple clusters to the storage system, connect the cluster nodes to the appropriate Gigabit or 10 Gigabit Ethernet switches and then connect the Gigabit or 10 Gigabit Ethernet switches to the appropriate storage processors on the processor enclosure.

For guidelines on iSCSI switch-attached clusters, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability Clustering website at [dell.com/ha](http://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 iSCSI Switch-Attached Clusters to the CX4-120, CX4-240, CX4-480 or CX4-960 Storage System***

- 1** In the first cluster, connect cluster node 1 to the iSCSI network:
  - a** Connect a local area network (LAN) cable from NIC-0 to switch 0 (sw0).
  - b** Connect a LAN cable from NIC-1 to switch 1 (sw1).
- 2** In the first cluster, repeat step 1 for each additional cluster node.
- 3** For each additional cluster, repeat step 1 and step 2.
- 4** Connect the storage system to the iSCSI network:
  - a** Connect a LAN cable from switch 0 (sw0) to SP-A iSCSI port 0.
  - b** Connect a LAN cable from switch 0 (sw0) to SP-B iSCSI port 0.
  - c** Connect a LAN cable from switch 1 (sw1) to SP-A iSCSI port 1.
  - d** Connect a LAN cable from switch 1 (sw1) to SP-B iSCSI port 1.



**NOTE:** Additional cables can be connected from the network switches to the storage system if there are available iSCSI ports on the storage processors.

### **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 iSCSI network. 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.


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 and 2008 operating systems support mount points, allowing greater than 22 drives per cluster.

Figure 2-13 provides an example of cabling the cluster nodes to four Dell/EMC storage systems.


**Figure 2-13. PowerEdge Cluster Nodes Cabled to Four Storage Systems**







# Preparing Your Systems for Clustering

 **WARNING:** Only trained service technicians are authorized to remove and access any of the components inside the system. See your *safety information* 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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

- 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" on page 61 and the iSCSI information in the "iSCSI Configuration Worksheet" on page 59.

- 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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

- 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 37.



**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 37.
- 10 Configure roles for cluster networks.
- 11 Test the failover capabilities of your new cluster.

- 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.

See the *Dell Cluster Configuration Support Matrix* on the Dell High Availability Clustering website at [dell.com/ha](http://dell.com/ha) for a list of:

- Dell PowerEdge Systems
- Gigabit or 10 Gigabit Ethernet network interface cards (NICs)
- Recommended list of operating system variants, specific driver and firmware revisions

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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

The following sub-sections describe steps that must be taken to enable communication between the cluster nodes and your shared Dell/EMC CX4-series iSCSI storage array, and to present disks from the storage array to the cluster.

## Installing the iSCSI NICs

It is recommended that you install the latest supported version of the driver. If the NIC driver requires any service packs or hotfixes to be installed along with the operating system, install them at this time. If supported by the network switches, enable jumbo frame to increase network payload size and performance.

For information about supported NICs and drivers, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at [dell.com/ha](http://dell.com/ha).

## Installing the Microsoft iSCSI Software Initiator

- 1 Use a web browser and go to the Microsoft Download Center website at [microsoft.com/downloads](http://microsoft.com/downloads).
- 2 Search for iSCSI initiator.
- 3 Select and download the latest supported initiator software and related documentation for your operating system.



**NOTE:** For the latest supported Software Initiator version, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at [dell.com/ha](http://dell.com/ha).

- 4 Double-click the executable file. The installation wizard launches.
- 5 In the Welcome screen, click Next.
- 6 In the following screens select the Initiator Service and Software Initiator. Click **Next** to continue with the installation.
- 7 Read and accept the license agreement and click **Next** to install the software.
- 8 At the completion screen, click Finish to complete the installation.
- 9 Select the **Do not restart now** option to reboot the system after modifying the TCP/IP registry settings in the section "Installing and Configuring the Shared Storage System" on page 41.



## Modifying the TCP Registry Settings

- 1 Determine the IP addresses or the DHCP IP addresses that are used for iSCSI traffic.
- 2 Start the Registry Editor.
  - a Select **Start**→**Run**.
  - b Type `Regedit` and click **OK**.
- 3 Locate and click the registry subkey located at the following location:  
**HKEY\_LOCAL\_MACHINE**→**SYSTEM**→**CurrentControlSet**→**Services**→**TCPIP**→**Parameters**→**Interfaces**
- 4 Click each interface GUID associated with the iSCSI networks, and perform the following steps:
  - a Select **Edit**→**New**→**DWORD** value.
  - b Name the new value **TcpAckFrequency**.
  - c Assign a value of 1.
- 5 Exit the Registry Editor.

## 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 Access Control (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 Control and Creating Storage Groups Using Navisphere" on page 45 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 Control

iSCSI 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 need to enable **Access Control**.

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 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 Control** 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**, 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 3-1 describes the properties in the storage group.

**Table 3-1. Storage Group Properties**

<b>Property</b>	<b>Description</b>
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"><li>• <b>Name</b> – Name of the host system</li><li>• <b>IP address</b> – IP address of the host system</li><li>• <b>OS</b> – Operating system that is running on the host system</li></ul> <b>NOTE:</b> In a clustered environment, all nodes of a cluster must be connected to the same storage group.
LUNs in storage group	Lists the LUNs in the storage group. Each LUN entry contains the following fields: <ul style="list-style-type: none"><li>• <b>Identifier</b> – LUN icon representing the LUN</li><li>• <b>Name</b> – Name of the LUN</li><li>• <b>Capacity</b> – Amount of allocated storage space on the LUN</li></ul>

## 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

In a cluster environment, a NIC may have multiple IP addresses bound to it. To ensure that the agent registers the host to the storage system with the correct IP address, perform the following steps:


- 1 Using a text editor to create a file named **agentID.txt** in the **C:\Program Files\EMC\Navisphere** directory (or **C:\Program Files (x86)\EMC\Navisphere** for Windows Server 2003 (64-bit) systems).
- 2 Add the following two separate lines to the **agentID.txt** file, with no special formatting:
  - First line: Fully qualified hostname. For example, enter **node1.domain1.com**, if the host name is *node1* and the domain name is *domain1*.
  - Second line: IP address that you want the agent to register and use to communicate with the storage system.

## EMC PowerPath

EMC PowerPath<sup>®</sup> automatically reroutes iSCSI traffic from the host system and a Dell/EMC CX4-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 Control and Creating Storage Groups Using Navisphere

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

 **CAUTION:** 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.

## 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.

## 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.

## Optional Storage Features

Your Dell/EMC CX4-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 Control. For specific version requirements, see the *Dell Cluster Configuration Support Matrix* on the Dell High Availability Clustering website at [dell.com/ha](http://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* available at [support.dell.com/manuals](http://support.dell.com/manuals).

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* available at [support.dell.com/manuals](http://support.dell.com/manuals).



# 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**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
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 17 for more information.
	The length of the interface cables exceeds the maximum allowable length.	
	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"> <li>• All switched zones are configured correctly.</li> <li>• The Access Control software is enabled on the storage system.</li> <li>• All LUNs and hosts are assigned to the proper storage groups.</li> </ul>
	The cluster is in a SAN, and one or more zones are not configured correctly.	Verify the following: <ul style="list-style-type: none"> <li>• Each zone contains only one initiator.</li> <li>• Each zone contains the correct initiator and the correct storage port(s).</li> </ul>

**Table A-1. General Cluster Troubleshooting (continued)**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
One of the nodes takes a long time to join the cluster.	The node-to-node network has failed due to a cabling or hardware failure.	Check the network cabling. Ensure that the node-to-node interconnection and the public network are connected to the correct NICs.
or		
One of the nodes fail to join 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 <a href="http://support.microsoft.com">support.microsoft.com</a> for more information.
	Long delays in node-to-node communications may be normal.	Verify that the nodes can communicate with each other by running the <b>ping</b> command from each node to the other node. Try both the host name and IP address when using the <b>ping</b> command.

**Table A-1. General Cluster Troubleshooting (continued)**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
Attempts to connect to a cluster using Cluster Administrator fail.	<p>The Cluster Service has not been started.</p> <p>A cluster has not been formed on the system.</p> <p>The system has just been booted and services are still starting.</p> <p>The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes.</p>	<p>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:</p> <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> <p>Configure the Internet Connection Firewall to allow communications that are required by MSCS and the clustered applications or services.</p> <p>See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at <a href="http://support.microsoft.com">support.microsoft.com</a> for more information.</p>

**Table A-1. General Cluster Troubleshooting (continued)**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
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" section 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> available at <a href="http://support.dell.com/manuals">support.dell.com/manuals</a> .
	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)**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
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"><li>• Check all cable connections</li><li>• Check all zone configurations</li><li>• Check the Access Control settings on the attached storage systems</li><li>• Use the <b>Advanced with Minimum</b> option</li></ul> <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 <a href="http://support.microsoft.com">support.microsoft.com</a> 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.</p> <p>If you are running Windows Server 2003, this situation is normal if the cluster node does not own the cluster disk.</p>	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.	<p>Perform the following steps:</p> <ol style="list-style-type: none"> <li><b>1</b> On the Windows desktop, right-click <b>My Computer</b> and click <b>Manage</b>.</li> <li><b>2</b> In the <b>Computer Management</b> window, double-click <b>Services</b>.</li> <li><b>3</b> In the <b>Services</b> window, double-click <b>Cluster Services</b>.</li> <li><b>4</b> In the <b>Cluster Services</b> window, click the <b>Recovery</b> tab.</li> <li><b>5</b> Click the <b>First Failure</b> drop-down arrow and select <b>Restart the Service</b>.</li> <li><b>6</b> Click the <b>Second Failure</b> drop-down arrow and select <b>Restart</b> the service.</li> <li><b>7</b> Click <b>OK</b>.</li> </ol> <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 <a href="http://support.microsoft.com">support.microsoft.com</a> and the Microsoft Technet website at <a href="http://microsoft.com">microsoft.com</a>.</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.	<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 <a href="http://support.microsoft.com">support.microsoft.com</a> for more information.</p>



**Table A-1. General Cluster Troubleshooting (continued)**

<b>Problem</b>	<b>Probable Cause</b>	<b>Corrective Action</b>
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 <a href="http://support.microsoft.com">support.microsoft.com</a> for more information.
The Storage Cluster Validation test fails while using the Persistent Reservation command.	PowerPath is installed before the networking has been configured. The host ID is duplicate among the cluster nodes as a result.	Run the <code>emphostid set -no_prompt</code> command on each cluster node to correct the host IDs. See the EMC PowerPath document for details on how to use <code>emphostid</code> utility.



## iSCSI Configuration Worksheet

*If you need additional space for more than one host server, use an additional sheet.*

<b>A</b>	Static IP address (host server)	Subnet	Default Gateway
Server 1, iSCSI NIC port 0	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
Server 1, iSCSI NIC port 1	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
Server 2, iSCSI NIC port 0	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
Server 2, iSCSI NIC port 1	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
Server 3, iSCSI NIC port 0	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
Server 3, iSCSI NIC port 1	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __

**Mutual CHAP  
Secret**

---

<b>B</b>	Static IP address (storage array)	Subnet	Default Gateway
SP-A, iSCSI port 0	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
SP-A, iSCSI port 1	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
SP-B, iSCSI port 0	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __
SP-B, iSCSI port 1	__ . __ . __ . __	__ . __ . __ . __	__ . __ . __ . __

**Target CHAP  
Secret**

---



# 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 C-1. Cluster Information**

<b>Cluster Information</b>	<b>Cluster Solution</b>
Cluster name and IP address	
Server type	
Installer	
Date installed	
Applications	
Location	
Notes	

**Table C-2. Cluster Node Information**

<b>Node Name</b>	<b>Service Tag Number</b>	<b>Public IP Address</b>	<b>Private IP Address</b>

**Table C-2. Cluster Node Information (continued)**

<b>Node Name</b>	<b>Service Tag Number</b>	<b>Public IP Address</b>	<b>Private IP Address</b>

**Additional Networks**


**Table C-3. Storage Array Information**

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

# Index

## A

Access Control  
about, 42

## C

cable configurations  
cluster interconnect, 19  
for client networks, 19  
for mouse, keyboard, and  
monitor, 15  
for power supplies, 15

cluster  
optional configurations, 12

cluster configurations  
connecting to multiple shared  
storage systems, 33  
connecting to one shared storage  
system, 12  
direct-attached, 12, 20

cluster storage  
requirements, 9

clustering  
overview, 7

## D

Dell/EMC CX4-series, 7  
Cabling a Multi-Node Cluster, 24

Cabling a Two-Node Cluster, 22  
Cabling Description, 29  
Cabling Multiple Clusters, 25  
Cabling Multiple SAN-Attached  
Clusters, 32  
Cabling Storage, 26  
CX4-120, 10  
CX4-240, 10  
CX4-480, 10  
CX4-960, 10  
Zoning, 33

direct-attached cluster  
about, 20

## K

keyboard  
cabling, 15

## L

LUNs  
assigning to hosts, 47  
configuring and managing, 47

## M

MirrorView  
about, 11

- monitor
  - cabling, 15
- mouse
  - cabling, 15
- MSCS
  - installing and configuring, 49

## **N**

- Navisphere Agent
  - about, 44
- Navisphere Manager
  - about, 10, 43
  - hardware view, 10
  - storage view, 10
- network adapters
  - cabling the private network, 18-19
  - cabling the public network, 19

## **O**

- operating system
  - Windows Server 2003, Enterprise Edition
    - installing, 39

## **P**

- power supplies
  - cabling, 15
- PowerPath
  - about, 45

- private network
  - cabling, 17, 19
  - hardware components, 19
  - hardware components and connections, 19
- public network
  - cabling, 17

## **R**

- RAID
  - configuring the RAID level, 47

## **S**

- SAN-Attached Cluster, 13
- SAN-attached cluster configurations, 12
- shared storage
  - assigning LUNs to hosts, 47
- SnapView
  - about, 11
- storage groups
  - about, 42
- storage management software
  - Access Control, 42
  - Navisphere Agent, 44
  - Navisphere Manager, 43
  - PowerPath, 45
- storage system
  - configuring and managing LUNs, 47
  - configuring the hard drives, 46



## **T**

troubleshooting

- connecting to a cluster, 53

- shared storage subsystem, 51

## **W**

warranty, 13

