

Dell Compellent Storage  
Center Fibre Channel 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 potential damage to hardware or loss of data if instructions are not followed.



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

---

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

© 2011 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™, PowerVault™, and Compellent™ are trademarks of Dell Inc. Microsoft®, Windows®, Windows Server®, and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other trademarks and trade names may be used in this publication 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 . . . . .	5
	<b>Cluster Solution</b> . . . . .	<b>5</b>
	<b>Cluster Hardware Requirements</b> . . . . .	<b>6</b>
	Cluster Nodes. . . . .	6
	Cluster Storage . . . . .	7
	<b>Other Documents You May Need</b> . . . . .	<b>9</b>
2	Cabling Your Cluster Hardware . . . . .	11
	<b>Cabling the Mouse, Keyboard, and Monitor</b> . . . . .	<b>11</b>
	<b>Cabling the Power Supplies</b> . . . . .	<b>11</b>
	<b>Cabling Your Cluster for Public and Private Networks</b> . . . . .	<b>12</b>
	Cabling the Public Network . . . . .	13
	Cabling the Private Network . . . . .	13
	NIC Teaming . . . . .	14
	<b>Cabling the Storage System</b> . . . . .	<b>15</b>
3	Preparing Your Systems for Clustering . . . . .	25
	<b>Cluster Configuration Overview</b> . . . . .	<b>25</b>
	<b>Installation Overview</b> . . . . .	<b>27</b>
	<b>Installing the Fibre Channel HBAs</b> . . . . .	<b>27</b>

Installing the Fibre Channel HBA Drivers. . . . .	28
<b>Implementing Zoning on a Fibre Channel Switched Fabric . . . . .</b>	<b>28</b>
<b>Installing and Configuring the Shared Storage System. . . . .</b>	<b>29</b>
Setting Up the Controllers . . . . .	30
Create a Server . . . . .	31
Create a Server Cluster. . . . .	32
Create a Volume for the Server Cluster. . . . .	33
Optional Storage Features . . . . .	34
<b>Installing and Configuring a Failover Cluster . . . . .</b>	<b>35</b>
A Troubleshooting . . . . .	37
B Zoning Configuration Form . . . . .	43
C Cluster Data Form . . . . .	45

# 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 systems in a multi-tier enterprise application architecture in your cluster.

This document provides information to configure your Dell Compellent Storage Center Fibre Channel storage array with one or more Failover Clusters. It provides information on 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 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* 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 Failover Cluster, see the *Dell Cluster Configuration Support Matrices* at [dell.com/ha](http://dell.com/ha).

## Cluster Solution

Your cluster implements a minimum of two nodes to a maximum of sixteen nodes and provides the following features:

- 8 Gbps and 4 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:

- 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.
RAM	At least 1 GB of memory.
Host Bus Adapter (HBA) ports	Two Fibre Channel HBA ports 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.  <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.

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

Component	Minimum Requirement
Internal disk controller	<p>One controller connected to at least two internal hard drives for each node. Use any supported RAID controller or disk controller.</p> <p>Two hard drives are required for mirroring (RAID 1) and at least three are required for disk striping with parity (RAID 5).</p> <p><b>NOTE:</b> It is highly recommended that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives.</p>



**NOTE:** For more information about supported systems, HBAs, and operating system variants, see the *Dell Cluster Configuration Support Matrices* 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 system	Compellent Storage Center with dual Series-30 or Series-40 storage center controllers in a clustered configuration
Disk enclosure	<ul style="list-style-type: none"> <li>• 11 Fibre Channel loops, each with up to seven enclosures</li> <li>• 10 SAS chains, each with up to eight enclosures and 96 hard drives</li> <li>• 11 SATA loops, each with up to five enclosures</li> </ul>
Hard disk	<p>At least three for RAID 10 (2 data and 1 spare) and at least six for RAID 5 (5 data and 1 spare)</p> <p><b>NOTE:</b> RAID 6 and Dual Mirrored RAID 10 are also supported. Do not use RAID 0 in the cluster.</p>
Fibre Channel Front-End I/O card	Must support N_Port ID Virtualization (NPIV) in order to support Virtual Port mode.

**Table 1-2. Cluster Storage Requirements (continued)**

Hardware Components	Requirement
Fibre Channel switch	At least two 8 Gbps Fibre Channel switches. The switches must support NPIV in order to support Virtual Port mode.
Multiple clusters and stand-alone systems	Can share a storage system. See "Installing and Configuring the Shared Storage System" on page 29.



**NOTE:** NPIV allows multiple port IDs to share a single physical port.



**NOTE:** Virtual Port mode allows the Storage Center to expand the number of available front-end ports by transferring data on all the ports. If a physical port fails, the virtual port can be moved to a different physical port within the fault domain. It is recommended that Virtual Port mode be used in the cluster environment.

The storage system in the cluster is centrally managed by one host system (also called a *management station*) running Compellent Storage Center software—a centralized storage management application used to configure the Compellent Storage Center.

The Compellent Storage Center System Manager performs the following functions:

- Provides a central management interface to create and manage Storage Center volumes, servers, disks, and users.
- Displays the status of hardware components.
- Enables local and remote backup and restore.
- Provides Phone Home technical support.
- Allows multiple users to have different levels of access privileges.



Optional software for the shared storage system includes:

- **Data Progression**—leverages cost and performance differences between storage tiers, allowing the maximum use of lower-cost drives for stored data, while maintaining high performance drives for frequently-accessed data.
- **Data Instant Replay**—A Replay is a point-in-time copy of one or more volumes. Once an initial Replay of a volume is taken, subsequent Replays preserve pointers to data that has changed since the previous Replay. This minimizes the amount of storage space required to preserve periodic copies of a volume.
- **Remote Instant Replay**—replicates volumes to a remote Storage Center. It offers two modes: asynchronous and synchronous.

## 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.**

- 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 2008 Installation and Troubleshooting Guide* provides information on deploying your cluster with Windows Server 2008 operating systems.
- The *Storage Center System Manager Setup Guide* describes how to set up a new Storage Center.
- The *Storage Center System Manager User Guide* provides instructions for using Storage Center System Manager to manage storage.
- The *Enterprise Manager User Guide* provides instructions for managing multiple Storage Centers.
- 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.
- Release notes, updates, 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.



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

# Cabling Your Cluster Hardware

## 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

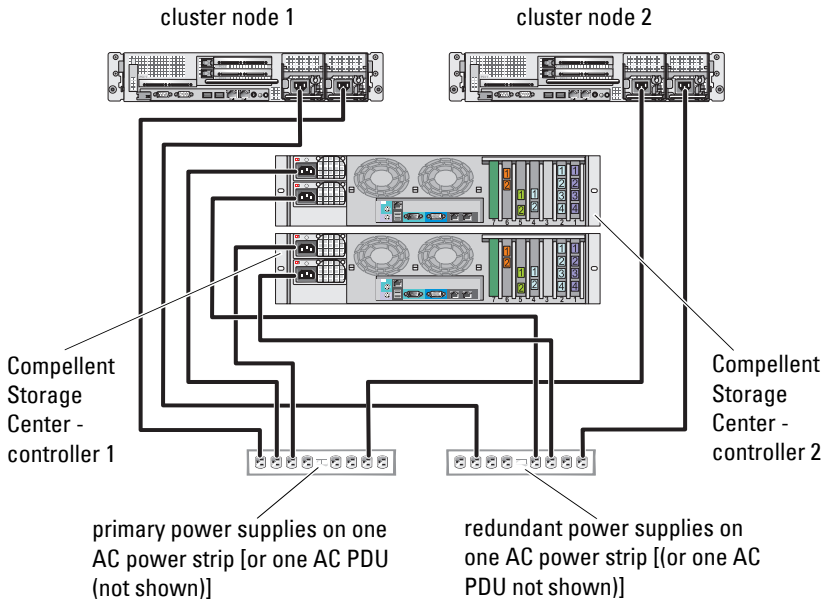
See the documentation for each component in your cluster solution and ensure that the specific power requirements are satisfied.

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 illustrates recommended methods for power cabling for a cluster solution consisting of two Dell PowerEdge systems and a dual controller Dell Compellent Storage Center. 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 Two Power Supplies in PowerEdge Systems**



## Cabling Your Cluster for Public and Private Networks


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

**Table 2-1. Network Connections**

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

Figure 2-2 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 2-2. 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	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 Connectors (LCs), or the twinax 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 standard CAT5e or CAT6 Ethernet cable between the network adapters in both nodes.
	Copper 10 Gigabit Ethernet network adapters with SFP+ connectors	Connect a twinax 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:** 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.



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

# Cabling the Storage System

This section provides information on cabling your cluster to a storage system in a SAN-attached configuration.

## Cabling a Cluster to a Compellent Storage Center Storage System

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

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

For more information on Fibre Channel switch fabrics, see "Implementing Zoning on a Fibre Channel Switched Fabric" on page 28.

Figure 2-3 shows an example of a two node SAN-attached cluster.


Figure 2-4 shows an example of an sixteen-node SAN-attached cluster.

Similar cabling concepts can be applied to clusters with 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 2-3. Two-Node SAN-Attached Cluster**





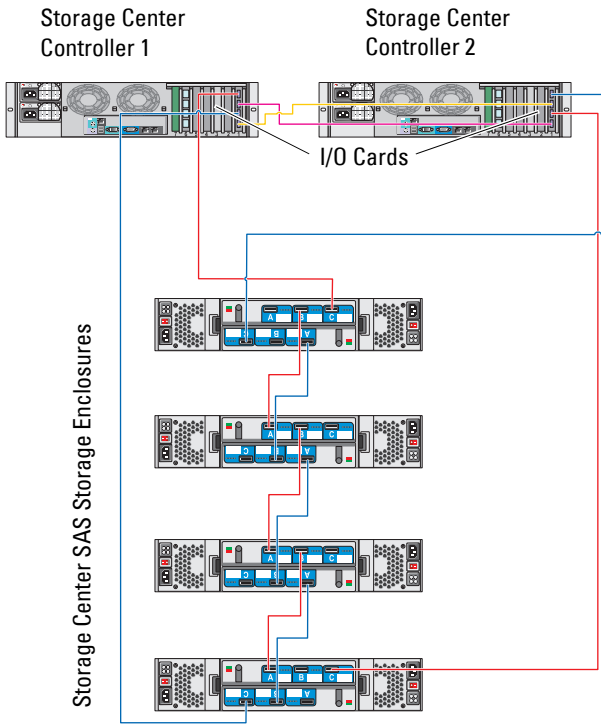
**Figure 2-4. Sixteen-Node SAN-Attached Cluster**



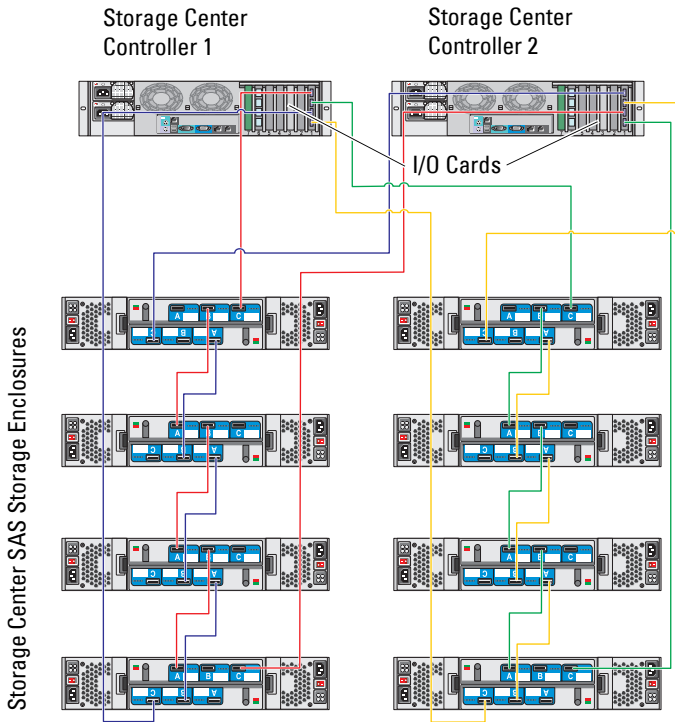
## Cabling the Compellent Storage Center Back-End

For information on how to cable the Compellent Storage Center back-end, see the Compellent Storage Center documentation. The following are two examples on how to connect the back-end cables.

**Figure 2-5. Back-End Cabling With One SAS Chain**



**Figure 2-6. Back-End Cabling With Multiple SAS Chains**



### **Cabling the Cluster Nodes and the Compellent Storage Center Front-End**

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 storage controller 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.

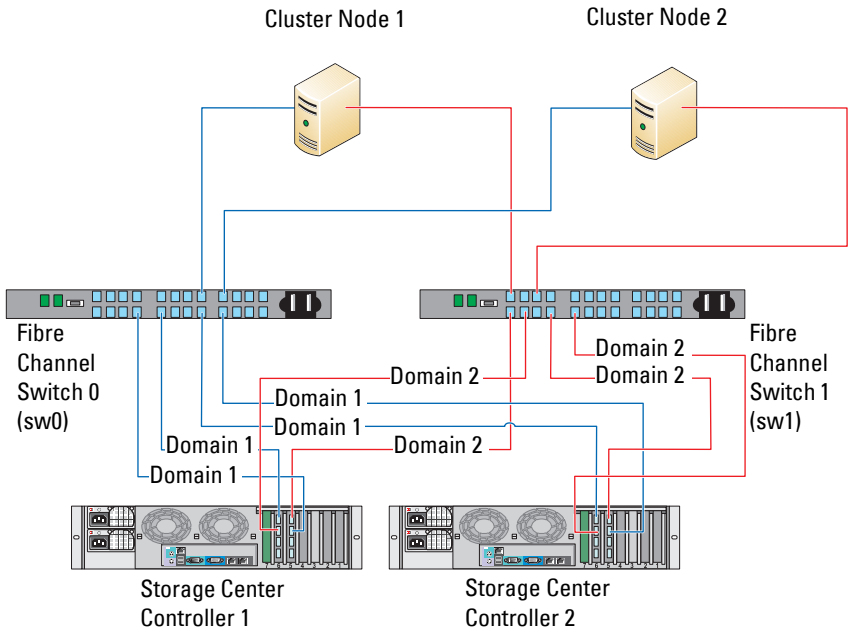
Each HBA port is cabled to a port on a Fibre Channel switch. One or more cables connect from the outgoing ports on a switch to a storage controller on a Compellent storage system.

- 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 additional cluster node.
- 3** Connect the storage system to the SAN:
  - a** Connect a cable from Fibre Channel switch 0 (sw0) to the first front-end fibre channel port on the Storage Center Controller 1.
  - b** Connect a cable from Fibre Channel switch 0 (sw0) to the first front-end fibre channel port on the Storage Center Controller 2.
  - c** Connect a cable from Fibre Channel switch 1 (sw1) to the second front-end fibre channel port on the Storage Center Controller 1.
  - d** Connect a cable from Fibre Channel switch 1 (sw1) to the second front-end fibre channel port on the Storage Center Controller 2.

Add more cables to support port failover and provide more bandwidth:

- e** Connect a cable from Fibre Channel switch 0 (sw0) to the third front-end fibre channel port on the Storage Center Controller 1.
- f** Connect a cable from Fibre Channel switch 0 (sw0) to the third front-end fibre channel port on the Storage Center Controller 2.
- g** Connect a cable from Fibre Channel switch 1 (sw1) to the fourth front-end fibre channel port on the Storage Center Controller 1.
- h** Connect a cable from Fibre Channel switch 1 (sw1) to the fourth front-end fibre channel port on the Storage Center Controller 2.

**Figure 2-7. Cabling a SAN-Attached Cluster to the Compellent Storage System**



**NOTE:** Additional cables can be connected from the fibre channel switches to the storage system if there are available front-end fibre channel ports on the storage processors.

### **Cabling Multiple SAN-Attached Clusters to a Compellent 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 the storage controllers.

#### ***Cabling Multiple SAN-Attached Clusters to the Compellent 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 additional cluster 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 the first front-end fibre channel port on the Storage Center Controller 1.
  - b Connect a cable from Fibre Channel switch 0 (sw0) to the first front-end fibre channel port on the Storage Center Controller 2.
  - c Connect a cable from Fibre Channel switch 1 (sw1) to the second front-end fibre channel port on the Storage Center Controller 1.
  - d Connect a cable from Fibre Channel switch 1 (sw1) to the second front-end fibre channel port on the Storage Center Controller 2.

Add more cables to support port failover and provide more bandwidth:

- e Connect a cable from Fibre Channel switch 0 (sw0) to the third front-end fibre channel port on the Storage Center Controller 1.
- f Connect a cable from Fibre Channel switch 0 (sw0) to the third front-end fibre channel port on the Storage Center Controller 2.
- g Connect a cable from Fibre Channel switch 1 (sw1) to the fourth front-end fibre channel port on the Storage Center Controller 1.
- h Connect a cable from Fibre Channel switch 1 (sw1) to the fourth front-end fibre channel port on the Storage Center Controller 2.

### 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 Fibre Channel switch.

Figure 2-8 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. 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 2-8. Cabling a Storage System and a Tape Library**



### **Obtaining More Information**

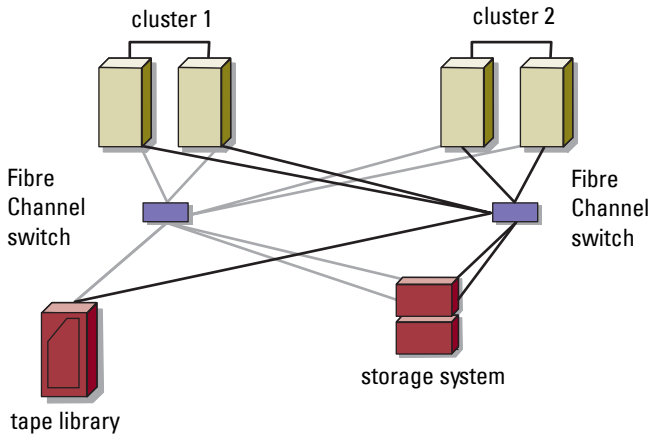
See the storage and tape backup documentation for more information on configuring these components.

### **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 2-9 provides an example of cabling the cluster nodes to your storage systems and SAN backup with a tape library.

**Figure 2-9. Cluster Configuration Using SAN-Based Backup**





# 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 through step 7 and step 10 through step 13, see the "Preparing your systems for clustering" section of *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* 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 (see "Cluster Data Form" on page 45), and the Zoning configuration information in the Zoning Configuration form (see "Zoning Configuration Form" on page 43).

- 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 Microsoft 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 2008 Installation and Troubleshooting Guide* 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 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. See "Preparing Your Systems for Clustering" on page 25.



**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. See "Preparing Your Systems for Clustering" on page 25.
- 10 Configure roles for cluster networks.
- 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 2008 R2, Enterprise x64 Edition. If the operating system varies among nodes, it is not possible to configure a Failover Cluster successfully. It is recommended that you 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 Matrices* at [dell.com/ha](http://dell.com/ha).

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

The following sub-sections describe the procedure to enable communication between the cluster nodes and your shared Dell Compellent 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 Matrices* at [dell.com/ha](http://dell.com/ha).

## Installing the Fibre Channel HBA Drivers

For more information about installing and configuring HBAs, see the following:

- Compellent HBAs—Compellent documentation that is included with your HBA kit.
- Emulex HBAs—Emulex support at [emulex.com](http://emulex.com) or the Dell Support at [support.dell.com](http://support.dell.com).
- QLogic HBAs—QLogic support at [qlogic.com](http://qlogic.com) or Dell Support at [support.dell.com](http://support.dell.com).

For more information about supported HBA controllers and drivers, see the *Dell Cluster Configuration Support Matrices* at [dell.com/ha](http://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 the Compellent Storage Center storage system in a switched fabric using Fibre Channel switch zoning and LUN mapping with Compellent Storage Center software. This configuration allows multiple clustered or nonclustered hosts to share a storage system.

PowerEdge cluster configurations support both World Wide Name (WWN) and port zoning.

A WWN is a unique numeric identifier assigned to Fibre Channel interfaces, such as HBA ports, storage controller 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. WWN port name zoning allows you to move cables between switch ports within the fabric without having to update the zones.

### Single Initiator Zoning

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

Follow these guidelines when you create single-initiator zones:

- Create a zone for each HBA port, all the Compellent front-end virtual ports, and the integrated bridge/SNC or fibre-channel interface on a tape library.
- Create a zone with only the Compellent physical WWNs.

## Installing and Configuring the Shared Storage System

This section describes how to install and configure the Compellent storage system in your cluster. The procedures, especially how to initialize or set up the controllers, provide a high-level description of the tasks. For more information, see the *Storage Center System Manager Setup Guide* and *Storage Center System Manager User Guide*. If you need support assistance, contact Dell Technical Support Services at [support.dell.com/compellent](http://support.dell.com/compellent).



**NOTE:** All hardware must be installed and cabled before you begin the setup and software installation process. The Fibre Channel switches must be configured and zoned before the controllers are set up.

The following pre-installation documentation is provided by your Storage Architect or Business Partner:

- List of hardware needed to support storage requirements
- Optional connectivity diagrams to illustrate cabling between the controllers, enclosures, network, and servers
- Optional network information, such as IP addresses, subnet masks, and gateways

These documents provide information about site-specific settings to configure the controllers.

## Setting Up the Controllers

- 1 Turn on each controller.
- 2 Connect to each controller.
- 3 Reset hardware and system serial numbers.
- 4 Set the IP addresses.
- 5 Configure the controllers using the Startup Wizard:
  - a Launch the **Storage Center Startup Wizard**.
  - b Accept the license agreement.
  - c Load the license.
  - d Create disk folder.
- 6 Add the second controller and perform the following settings:
  - Time settings
  - System setup
  - Configure SMTP
  - Update setup
  - User setup
  - Configure local ports
  - Generate SSL certificate
- 7 Complete the Storage Center setup:
  - a Configure a Phone Home proxy
  - b Phone Home

- c Check for Storage Center updates

## **Create a Server**

- 1 From the system tree in the **Storage Management** window, select the **Servers** node.
- 2 From the shortcut menu, select **Create Server**. The **Create Server Wizard** appears. The wizard lists Host Bus Adapters (HBAs) recognized by the Storage Center.
- 3 Select one or more HBAs belonging to the server.
- 4 Click **Continue**. A window allowing you to name the server is displayed.
- 5 Enter a name for the server or accept the default. Enter a folder name and any note (up to 255 characters).
- 6 From the drop-down menu, select **Windows 2008 MPIO**.  
Volumes are mapped to servers according to the rules of the server operating system. To view the rules of an operating system:
  - a Select a system.
  - b Click **Continue**.  
The screen displays the name and attributes of the system.
  - c Click on the question mark icon.  
A window opens displaying the rules of the selected operating system.
- 7 Click **Create Now**.
- 8 On the next window, choose from the following options:
  - Map the server to a Volume
  - Map a different Server to a Volume
  - Create Volume
  - Create Server
- 9 Close the **Create Server Wizard**.

## Create a Server Cluster

A server cluster is a collection of servers. A server that is a member of a server cluster is referred to as a *cluster node*. Volumes can be mapped directly to a server cluster. All volumes mapped to a server cluster are automatically mapped to all nodes in the cluster.

- 1 From the system tree in the **Storage Management** window, select the **Servers** node.
- 2 From the shortcut menu, select **Create Server Cluster**.
- 3 Choose **Add Existing Server**.
- 4 Expand a server folder to view servers within the folder, if necessary, and select a server.
- 5 Click **Continue**.

The **Create Server Cluster** window is displayed.

- 6 Repeat step 3 through step 5 until all the servers are added to the Server Cluster.
- 7 Once the **Create Server Cluster** window displays all the servers for this server cluster, click **Continue**.

A window allowing you to name the server cluster is displayed.

- 8 Name the server cluster or accept the default. Add notes if necessary.
- 9 Click **Continue**.

If volumes were mapped to individual server nodes, the system displays a list of mapped volumes and asks you to select volumes to map to the server cluster.



**NOTE:** The following are the default system selections:

- except boot volumes, the system selects all volumes mapped to the server nodes
- boot volumes are not selected to be mapped to the server cluster

- 10 Click **Continue**.
- 11 Click **Create Now**.
- 12 On the next window, map volumes or click **Close**.

The server cluster is displayed in the system tree.



## Create a Volume for the Server Cluster

Volumes are configured through the Configure Volume Defaults window. To create a volume for your cluster:

- 1 On the **Storage Management** window, select **Create→ Volume**.

The Create Volume Wizard asks you to enter a volume size.

- 2 Enter a volume size in GB, TB, or PB. The maximum size of a volume is 10 PB.



**NOTE:** If your User Volume Defaults allow you to modify cache settings or Storage Profiles, an **Advanced** button appears.

- 3 Click **Continue**.

If Data Instant Replay is licensed for your system, the **Replay Profile** window is displayed.

- 4 Select an existing Replay Profile or click **Create a New Replay Profile**.

- 5 Click **Continue**.

A window is displayed where you can select or create a folder and name the volume.

- 6 Select a folder where you want to create the volume or create a new folder.

- 7 Enter a volume name or accept the default.

- 8 Click **Continue**.

The system displays the attributes of the volume.

- 9 Review the attributes. The attributes depend on the available options when you create the volumes.

- 10 Click **Create Now**.

The next window presents options that are available after creating a volume.

Volumes are automatically mapped to the last server you selected. If the last volume was mapped to the same Cluster Server, select **Map this Volume to this Server**. Otherwise, select **Map this Volume to a Different Server**.

- 11 Select a cluster server to map to the selected volume.

- 12 Click **Continue**.

**13** Click **Create Now**.

The volume is mapped.

**14** If you want to create another volume for the cluster server, select **Create Volume** and repeat step 2 through step 14. Otherwise, click **Close** to exit the Wizard.

## Optional Storage Features

Your Compellent Storage Center storage array may be configured to provide optional features that can be used in conjunction with your cluster.

- **Enterprise Manager:** It is a separately-licensed application that manages and monitors multiple Storage Center systems. It greatly simplifies Remote Instant Replay.
- **Data Progression:** It leverages cost and performance differences between storage tiers, allowing the maximum use of lower-cost drives for stored data, while maintaining high performance drives for frequently-accessed data.
- **Data Instant Replay:** A Replay is a point-in-time copy of one or more volumes. Once an initial Replay of a volume is taken, subsequent Replays preserve pointers to data that has changed since the previous Replay. This minimizes the amount of storage space required to preserve periodic copies of a volume.
- **Remote Instant Replay:** Remote Instant Replay replicates volumes to a remote Storage Center and is an integral part of disaster recovery plan. It offers two modes: synchronous and asynchronous.
  - **Synchronous replication:** It makes sure that a write is successfully written to the remote system before returning a Successful Completion command to the server IO request. The Storage Center does not acknowledge completion of the write-back to the server until both the write IO to the local volume and the IO sent to the remote system are complete. This indicates that both the replicating volume and the replicated volume are fully synchronized and there is no data loss in the event of a failure on the source system.
  - **Asynchronous replication:** Asynchronous replication acknowledges a write IO back to the server as soon as it is completed on the source system. The write IO is also queued for delivery to the Remote system.

This allows for more efficient link utilization and data transfer optimization. It also means that in the event of a local failure, writes present on the source system may not be present on the remote system.

Remote Instant Replay can be initiated through either Storage Center or Enterprise Manager.

## **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 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* at [support.dell.com/manuals](http://support.dell.com/manuals).



# Troubleshooting

This appendix provides troubleshooting information for your cluster configuration. Table 4-1 describes general cluster problems you may encounter and the probable causes and solutions for each problem.

**Table 4-1. 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. For more information, see "Cabling Your Cluster for Public and Private Networks" on page 12.
	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.
	Volumes are not mapped correctly to the server nodes.	Verify the following: <ul style="list-style-type: none"> <li>• All switched zones are configured correctly.</li> <li>• The volumes are mapped correctly to the server nodes.</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 (Fibre Channel daughter card).</li> <li>• Each zone contains the correct initiator and the correct storage port(s).</li> </ul>

**Table 4-1. 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. or One of the nodes fails 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. For more information, see the Microsoft Knowledge Base article KB883398 at <a href="http://support.microsoft.com">support.microsoft.com</a> .  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 4-1. 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 the cluster is 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. For more information, see the Microsoft Knowledge Base article KB883398 at <a href="http://support.microsoft.com">support.microsoft.com</a>.</p>

**Table 4-1. 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" of the <i>Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide</i> 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.
Unable to add a node to the cluster.	The new node cannot access the shared disks.	Ensure that the new cluster node can enumerate the cluster disks using Windows Disk Administration. If the disks do not appear in Disk Administration, do the following: <b>1</b> Check all cable connections. <b>2</b> Check all zone configurations. <b>3</b> Use the <b>Advanced with Minimum</b> option.
	The shared disks are enumerated by the operating system differently on the cluster nodes.	
	One or more nodes may have the Internet Connection Firewall enabled, blocking RPC communications between the nodes.	



**Table 4-1. 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. For more information, see the Microsoft Knowledge Base article KB883398 at <a href="http://support.microsoft.com">support.microsoft.com</a> .



## Zoning Configuration Form

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



# 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 6-1. Cluster Information**

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

**Table 6-2. Cluster Node Information**

Node Name	Service Tag Number	Public IP Address	Private IP Address

<b>Additional Networks</b>

**Table 6-3. Storage Array Information**

<b>Controller</b>	<b>Service Tag Number</b>	<b>Management IP (eth0)</b>	<b>Private IP (eth1)</b>	<b>Controller Cluster IP</b>
1				
2				

# Index

## C

- cable configurations
  - cluster interconnect, 13
  - for client networks, 13
  - for mouse, keyboard, and monitor, 11
  - for power supplies, 11
- cabling, 21
  - multiple SAN-Attached clusters to a Compellent storage system, 21
- cluster storage requirements, 7
- clustering
  - overview, 5

## D

- drivers
  - installing and configuring Emulex, 28

## E

- Emulex HBAs
  - installing and configuring, 28
  - installing and configuring drivers, 28

## H

- HBA drivers
  - installing and configuring, 28
- host bus adapter
  - configuring the Fibre Channel HBA, 27

## K

- keyboard
  - cabling, 11

## M

- monitor
  - cabling, 11
- mouse
  - cabling, 11
- MSCS
  - installing and configuring, 35

## N

- network adapters
  - cabling the private network, 13
  - cabling the public network, 13

## **O**

overview  
installation, 27

## **P**

power supplies  
cabling, 11

private network  
cabling, 12-13  
hardware components, 14  
hardware components and  
connections, 13

public network  
cabling, 12

## **S**

SAN  
configuring SAN backup in your  
cluster, 23

single initiator zoning  
about, 29

## **T**

tape library  
connecting to a PowerEdge  
cluster, 22

troubleshooting  
connecting to a cluster, 39  
shared storage subsystem, 37

## **W**

warranty, 9

## **Z**

zones  
implementing on a Fibre Channel  
switched fabric, 28