

Dell™ PowerEdge™ Cluster Systems

**Using Dell Blade Servers in
a Dell PowerEdge High
Availability Cluster**

Notes, Notices, and Cautions



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



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



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

Information in this document is subject to change without notice.

© 2007 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, *PowerVault*, and *OpenManage* are trademarks of Dell Inc.; *Intel*, *Pentium*, and *Celeron* are registered trademarks of Intel Corporation; *Microsoft*, *Windows*, and *Windows Server* 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 document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

December 2007

Rev. A00

Contents

1	Introduction	5
	Overview	5
	Supported PowerEdge Blade Server Cluster Components	6
	PowerEdge Server Enclosure	6
	Supported Cluster Configurations	16
	Direct-Attached Cluster	16
	Network-Attached Cluster	17
	Blade Server Requirements	19
	Cluster Nodes	20
	Cluster Storage	21
	Other Documents You May Need	22
2	Cabling Your Blade Cluster Hardware	25
	Cabling Your Cluster For Public and Private Networks	25
	Cabling the Private Network	27
	Cabling the Public Network	28

Cabling the Storage Systems	29
Direct-Attached Cluster.	29
Network-Attached Cluster	32
Installing the Operating System and Configuring Your Cluster.	43
Maintaining Your Cluster	43

Introduction

This document provides information for installing and managing your Dell™ PowerEdge™ blade server in a Dell PowerEdge cluster system and specific information about clustering your PowerEdge blade server modules with supported Dell PowerVault™ and Dell|EMC storage systems. Use this document in conjunction with the *Installation and Troubleshooting Guide* for your supported Dell PowerEdge cluster solution. This document is intended for experienced IT professionals who need to configure the cluster solution, and for trained service technicians who perform upgrade and maintenance procedures.

Overview

Clustering uses specific hardware and software to join multiple systems together to function as a single system and provide an automatic failover solution. If one of the cluster nodes (also referred to as nodes) fails, resources running on the failed system are moved (or failed over) to one or more systems in the cluster by the Microsoft® Cluster Service (MSCS) software. MSCS is the failover software component in specific versions of the Microsoft Windows® operating system. Cluster nodes share access to external storage systems; however, only one of the nodes can own any virtual disk or Logical Unit Number (LUN) in the external storage system at any time. MSCS controls which node has access to each virtual disk in the shared storage system.



NOTE: This document uses the terms *cluster node* and *server module* interchangeably.

When the failed system is repaired and brought back online, resources automatically transfer back (or fail back) to the repaired system or remain on the failover system, depending on how MSCS is configured. For more information about MSCS, see the *Installation and Troubleshooting Guide*.



NOTE: Reference to Microsoft Windows Server® 2003 in this guide implies reference to Windows Server 2003 Enterprise Edition, Windows Server 2003 R2 Enterprise Edition, Windows Server 2003 Enterprise x64 Edition, and Windows Server 2003 R2 Enterprise x64 Edition unless explicitly stated.

Supported PowerEdge Blade Server Cluster Components

The following sections discuss the various cluster components that are supported with your PowerEdge blade server cluster configuration.


PowerEdge Server Enclosure

The Dell PowerEdge cluster solution supports the Dell PowerEdge blade server enclosures. These systems function as enclosures for multiple PowerEdge server modules that can be configured together into multiple clusters. The system is configured with internal connections and hot-pluggable hardware components that provide the necessary communication links between the individual server modules (private network), the client network (public network), and an additional PowerEdge server enclosure.


Multiple PowerEdge server enclosures can be connected together in a cluster configuration. For supported cluster configurations, see "Cabling Your Blade Cluster Hardware" on page 25.

Figure 1-1 provides a front and back view of the PowerEdge 1855/1955 system. Figure 1-2 provides a front and back view of the PowerEdge M1000e system.

Figure 1-1. PowerEdge 1855/1955 Server Enclosure Overview




front view




back view

- | | | | |
|---|--|----|--|
| 1 | server module (10) | 8 | power supply module (4) |
| 2 | front-panel indicators | 9 | blanks (2) |
| 3 | Ethernet switch module or Ethernet pass-through module (not shown) (2) | 10 | Ethernet switch module (2) or Ethernet pass-through module (not shown) (2) |
| 4 | I/O bay 1 | 11 | I/O bay 4 |
| 5 | I/O bay 3 | 12 | I/O bay 2 |
| 6 | KVM module | 13 | fan modules (2) |
| 7 | DRAC/MC module | | |

Figure 1-2. PowerEdge M1000e Server Enclosure Overview



front view



back view

1	server module	8	I/O bay A2
2	power switch and KVM ports	9	I/O bay B2
3	control panel	10	I/O bay C2
4	DRAC/MC module	11	I/O bay C1
5	KVM module	12	I/O bay B1
6	fan module	13	I/O bay A1
7	power supply module		

The power supply modules, fan modules, Dell Remote Access Controller/Modular Chassis (DRAC/MC), and I/O modules are shared resources of the server modules in the chassis. The system may ship with an optional external Universal Serial Bus (USB) diskette drive and an optional external USB CD drive, which you can use to set up and configure the server modules.

For information on supported cluster configurations, see "Cabling Your Blade Cluster Hardware" on page 25. For a list of supported hardware and software components, see the *Platform Guide* for your PowerEdge cluster solution located on the Dell Support website at support.dell.com.



NOTE: To ensure proper operation and cooling, all bays must be populated with either a server module or a blank prior to turning on the system.

The PowerEdge server enclosure include the following hardware components for a cluster configuration:

- Server modules
- Gigabit Ethernet switch modules or Gigabit Ethernet pass-through modules (based on your configuration)
- Fibre Channel switch modules or Fibre Channel pass-through modules (based on your configuration)

Dell Remote Access Controller/Modular Chassis (DRAC/MC)

The DRAC/MC is a management module located in the back of the blade server system chassis that provides all of the chassis management functionality. The DRAC/MC provides serial and out-of-band Ethernet management ports to allow for management of the chassis and some basic blade functions.

The following is a list of features available on the DRAC/MC. Your system may have updates that enable additional features. Refer to the latest Dell Remote Access Controller/Modular Chassis User's Guide on the Dell Support website at support.dell.com.

- Remote management and monitoring of a system through the DRAC/MC Web-based graphical user interface (GUI), serial connection, or telnet connection.
- Access to the chassis System Event Log (SEL) and DRAC/MC logs.
- Integrated launch of the DRAC/MC interface from the Dell OpenManage™ IT Assistant.
- Ability to alert you to potential problems on the DRAC/MC by sending either an e-mail message or an SNMP trap through the DRAC/MC NIC to a management station.
- Ability to configure the DRAC/MC and update DRAC/MC firmware using a telnet session, a web-based user interface, or through a terminal session (for example, a hyperterminal or similar program).
- Ability to manage controller configurations, I/O modules configurations and settings, and perform power management functions such as shutdown, power up, and reset, from a telnet session.
- Web-based interface password-level security management.
- Role-based authority that provides assignable permissions for different systems management tasks.

Server Modules

Depending on the PowerEdge blade server module you choose for your cluster, the number of available expansion ports and dual inline memory modules (DIMMs) varies. Table 1-1 provides details about the supported PowerEdge server modules and PowerEdge server enclosures.

Table 1-1. PowerEdge Blade Server Features

Server Enclosure	1855/1955	M1000e
Supported blade server modules	PowerEdge 1855, 1955	PowerEdge M600, M605
Maximum number of server modules per server enclosure	10	16
Number of mezzanine card slots per server module	1	2

In a Dell PowerEdge high availability cluster configuration, each server module requires at least one expansion card. Expansion cards for either Fibre Channel or Gigabit Ethernet (for iSCSI) are available, and allow the server module to communicate with the shared storage system for the cluster. The expansion cards, also known as daughter cards or mezzanine cards, are installed on the server module and contain two I/O ports. These I/O ports are internally connected to two separate I/O modules in the server enclosure. By attaching to two separate I/O modules, an expansion card can provide redundant paths and load balance the I/O from the server module to the shared storage system(s) for the PowerEdge cluster solution. Table 1-1 outlines the number of expansion cards that are available on each supported server module.

Each server module is also configured with two additional Ethernet ports for cluster interconnects. These are internally connected to two separate Gigabit Ethernet pass-through or Gigabit Ethernet switch modules in the server enclosure. With certain server modules, it is also possible to configure additional Ethernet ports, which can enable the use of NIC teaming on the cluster public network.

For a list of supported Fibre Channel and iSCSI daughter cards, see the *Platform Guide* for your PowerEdge cluster located on the Dell Support website at support.dell.com.

Gigabit Ethernet

Gigabit Ethernet provides the following network connections, depending on your configuration:

- Public connection from the server modules to the public network
- Private node-to-node connection between the server modules for the private network
- iSCSI connection between the server modules and storage system(s). For more information, see "Supported Cluster Configurations" on page 16.

The server modules include two integrated Gigabit Ethernet network interface cards (NICs). You must configure at least two networks for each PowerEdge cluster solution. One cluster network is configured for heartbeat communications (private network) and is only accessible to the server modules in the cluster. The other cluster network is configured for the client network (public network) and is accessible to client systems. It is recommended that you configure the same network interface on each node for the same role in the cluster. For example, you can use the two integrated NICs to provide the private and public cluster networks. Such a configuration allows an expansion card in the server module to provide the I/O interface for the shared storage system(s) in the cluster. Use a Gigabit Ethernet expansion card for iSCSI shared storage systems and use a Fibre Channel expansion card for Fibre Channel shared storage systems. For more information about iSCSI and Fibre Channel clusters, see "Supported Cluster Configurations" on page 16.

Gigabit Ethernet Switch Module

The Gigabit Ethernet switch module provides a switched connection to the integrated NICs on each server module. Using the internal connections in the system chassis, the Gigabit Ethernet switch module can be used to provide the following configurations:

- A switched connection to the client network (public network).
- Network connection to one or more server modules in the Cluster configuration (private network).
- iSCSI connection between the server modules and storage system(s). For more information, see "Supported Cluster Configurations" on page 16.

Gigabit Ethernet Pass-Through Module

The Gigabit Ethernet pass-through module provides a non-switched connection between the server modules and an external Gigabit Ethernet device.



 **NOTE:** The Gigabit Ethernet pass-through module ports are preset to communicate at 1000 Mb and will not auto-negotiate to a slower speed. As a result, the Gigabit Ethernet pass-through module only supports a 1000-Mb connection.

Table 1-2 summarizes the supported Ethernet module configurations.

Table 1-2. Supported Ethernet Module Configurations

Ethernet Switch Module for iSCSI	Ethernet Pass-Through Module for iSCSI
Switch-attached configuration to four supported Dell EMC iSCSI storage systems	Direct-attached configuration to a Dell EMC iSCSI storage system
Switch-attached configuration to an external iSCSI network with up to four supported Dell EMC iSCSI storage systems	Switch-attached configuration to an external iSCSI network with up to four supported Dell EMC iSCSI storage systems

 **NOTE:** The supported iSCSI cluster configurations cable connections are similar to that of the Fibre Channel Cluster configuration. Each server module connects to the storage system using CAT5e or CAT6 LAN cables with RJ45 connectors that attach to Gigabit Ethernet daughter cards in the server modules and the Gigabit iSCSI ports on the cluster shared storage system.

Fibre Channel Module

You can configure the PowerEdge blade cluster with two hot-pluggable Fibre Channel switch modules installed on a PowerEdge 1855/1955 system or up to four hot-pluggable Fibre Channel switch/pass-through modules installed on the PowerEdge M1000e system to provide Fibre Channel connection between the server modules and storage system(s).

For more information about the Fibre Channel modules, see your Dell PowerEdge system documentation.

Fibre Channel Switch Module

The Fibre Channel switch module provides a *switched* connection between the Fibre Channel daughter card in the server modules and a supported Fibre Channel device. The switch module functions as a director, mapping requests and responses between the interconnected devices.

Additionally, the Fibre Channel switch module includes an internal serial port that communicates with the DRAC/MC module.

The Fibre Channel switch module supports the following configurations:

- Network-attached configuration with up to two supported Dell|EMC storage systems
- Network-attached connection to an external storage area network (SAN) with up to four supported Dell|EMC storage systems

Fibre Channel Pass-Through Module

The Fibre Channel pass-through module provides a *direct* connection between the Fibre Channel daughter card in the server modules and a supported Fibre Channel device. The pass-through module functions as a dedicated link between the server modules and specific ports on the pass-through module.

The pass-through module supports the following configurations:

- Direct-attached configuration to a supported Dell|EMC storage system
- Network-attached connection to an external SAN with up to four supported Dell|EMC storage systems

Table 1-3 summarizes the supported Fibre Channel module configurations.

Table 1-3. Supported Fibre Channel Module Configurations

Fibre Channel Switch Module	Fibre Channel Pass-through Module
Network-attached configuration to one or two supported Dell EMC Fibre Channel storage systems using embedded Fibre Channel switch modules	Direct-attached configuration to a Dell EMC Fibre Channel storage system
Network-attached configuration to an external SAN with up to four supported Dell EMC Fibre Channel storage systems	Network-attached connection to an external SAN with up to four supported Dell EMC Fibre Channel storage systems

Supported Cluster Configurations


The PowerEdge blade servers support both Fibre Channel and iSCSI cluster configurations using either Fibre Channel mezzanine cards or Gigabit Ethernet mezzanine cards to connect to the corresponding shared storage system in either direct-attached or network-attached environments.

Direct-Attached Cluster

In a direct-attached cluster, the cluster nodes are directly attached to a single storage system. The HBAs (mezzanine cards/daughter cards) in the nodes are internally connected to the pass-through modules. The pass-through modules are connected by cables directly to the RAID controllers (or storage processors) on the storage system.

Figure 1-3 shows the logical components of a direct-attached cluster configuration.

Figure 1-3. Direct-Attached Cluster



Network-Attached Cluster

Network-attached clusters provide configuration flexibility, expandability, and performance. In a network-attached cluster, all of the nodes (server modules) are attached to redundant switch fabrics. A network-attached cluster supports the following configurations:

- One or two storage systems using embedded switch fabrics
- Up to four storage systems using external fabrics

Figure 1-4 shows the network-attached cluster connected to an external SAN using embedded switch modules.

Figure 1-5 shows the network-attached cluster connected to an external SAN using embedded pass-through modules.

Figure 1-6 shows the network-attached cluster connected to an external SAN using embedded switch modules


 **NOTE:** Figure 1-4 through Figure 1-6 are for illustration only. Some cluster connections shown below are routed internally through the PowerEdge server enclosure.

Figure 1-4. Network-Attached Cluster Connected to an External SAN Using Embedded Switch Modules




Figure 1-5. Network-Attached Cluster Connected to an External SAN Using Embedded Pass-Through Modules

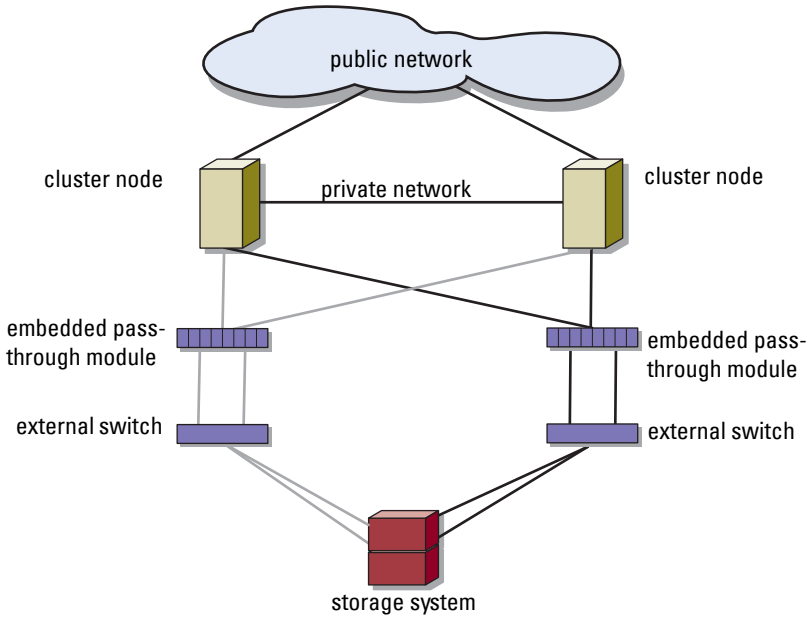



Figure 1-6. Network-Attached Cluster Connected to an External SAN Using Embedded Switch Modules



NOTE: It is recommended that you configure the paired inter-switch links (ISLs) as shown in Figure 1-6 for clusters with Fibre Channel storage systems. In an iSCSI cluster configuration, multiple connections between internal switch modules and external switches may not be practical depending on the Ethernet switch features.

Blade Server Requirements

Your cluster requires the following components:

- Server modules (cluster nodes)
- Fibre Channel or iSCSI storage subsystem interconnects
- Cluster interconnects (private network)
- Client network connections (public network)
- Operating system and storage management software

Cluster Nodes

Table 1-4 lists the hardware requirements for the server modules in the PowerEdge server enclosure.

Table 1-4. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	Two to eight PowerEdge server modules running the appropriate Windows Server 2003 operating system edition as described in Table 1-1
RAM	At least 512 MB of RAM installed per server module
Fibre Channel mezzanine card	If the cluster is configured with a Fibre Channel shared storage system, at least one dual-port Fibre Channel mezzanine card per server module
iSCSI mezzanine card	If the cluster is configured with an iSCSI shared storage system, at least one dual-port Gigabit Ethernet mezzanine card per server module
Gigabit Ethernet (cluster networks)	At least two Gigabit Ethernet interfaces per server module
RAID controller (optional)	One controller connected to two internal hard drives for each server module in a RAID 1 configuration NOTE: It is strongly recommended that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives.

Cluster Storage

Table 1-5 lists supported storage systems and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems.

Table 1-5. Cluster Storage Requirements

Storage Configuration	Requirement
Supported storage	See the <i>Platform Guide</i> for your PowerEdge cluster solution on the Dell Support website at support.dell.com .
Cluster shared storage	<p>If you are using either switch modules or pass-through modules to connect to network storage devices that are attached to external switches in a SAN or an IP SAN, see the <i>Platform Guide</i> for your PowerEdge cluster solution located on the Dell Support website at support.dell.com to determine the maximum number of supported storage systems.</p> <p>If you are using switch modules to connect to network storage devices in an internal SAN or IP SAN without external switches, you can use a maximum of two supported storage systems.</p> <p>If you are using pass-through modules to connect to storage devices in a direct-attached configuration, you can use a single supported storage system.</p>
Multiple clusters and stand-alone systems	Network storage systems that provide LUN masking or assignment functionality may be shared between multiple clusters or between clustered and stand-alone servers. To determine whether the configuration is supported, see the <i>Platform Guide</i> for your PowerEdge cluster solution located on the Dell Support website at support.dell.com .

Other Documents You May Need



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

- The *Platform Guide* provides information about the platforms that support your cluster configuration.
- The *Rack Installation Guide* included with your rack solution describes how to install your system into a rack.
- The *Getting Started Guide* provides an overview of initially setting up your system.
- The Dell PowerEdge system documentation provides information about system features, technical specifications, describes how to troubleshoot the PowerEdge server enclosure, and install or replace system components.
- The *Dell Remote Access Controller/Modular Chassis User's Guide* provides detailed information about using the remote management features of the system.
- The *Dell PowerConnect 5316M User's Guide* describes the Ethernet switch module features and basic operation.
- 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.
- The *Dell PowerEdge Expandable RAID Controller 4/iM (PERC 4/iM) Integrated Mirroring Guide* describes the PERC 4/iM integrated mirroring features.
- The *Dell™ PowerEdge™ Expandable RAID Controller 5/i and 5/E User's Guide* describes installation, configuration and troubleshooting procedures for PERC 5/i and PERC 5/E.
- Documentation for any components you purchased separately provides information to configure and install these options.
- The Dell PowerVault™ tape library documentation provides information for installing, troubleshooting, and upgrading the tape library.
- The documentation that came with your storage system.


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



NOTE: Always check for updates on the Dell Support website at support.dell.com and 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 Blade Cluster Hardware

 **NOTE:** The cluster nodes represent the individual server modules in your Dell™ PowerEdge™ server enclosure. The designations namely, node 1 and node 2 and the server enclosure are used for illustrative purposes only. The system and the storage components in your cluster configuration may vary.

Cabling Your Cluster For Public and Private Networks


The network interface cards (NICs) in the cluster nodes provide at least two network connections for each node, as described in Table 2-1.

Table 2-1. Network Connections

Network Connection	Description
Public network	<ul style="list-style-type: none"> • All connections to the client LAN. • The public network must be configured for Mixed mode for private network failover.
Private network	A dedicated connection for sharing cluster health and status information only.

Figure 2-1 shows an example of network cabling for a PowerEdge M1000e server enclosure in which dedicated NICs in each server module are internally connected to each other through a Gigabit Ethernet switch module (for the private network) and the remaining NICs are connected to the public network

Figure 2-1. Network Cabling Connection for PowerEdge M1000e Server Enclosure



- | | | | |
|---|---|---|--|
| 1 | public network Ethernet switch or pass-through module | 6 | node 2 |
| 2 | public NIC port | 7 | private NIC port |
| 3 | private NIC port | 8 | public NIC port |
| 4 | node 1 | 9 | private network Ethernet switch or pass-through module |
| 5 | internal connections | | |

Cabling the Private Network

The private (heartbeat) network is used for sharing cluster health and status information between the cluster nodes. Each cluster node connects to the private network through the second embedded NIC port on the server module system board. The second NIC port connects to the internal connections in the system chassis to the second Ethernet switch module or the associated port in the corresponding Gigabit pass-through module. The Gigabit Ethernet switch module or Gigabit pass-through module is connected internally to the private network-configured Gigabit Ethernet switch module in the private network or the corresponding port on the corresponding Gigabit Ethernet pass-through module.



NOTE: It is recommended that you configure the second NIC port on each server module for the private network.

The Gigabit Ethernet pass-through module connectors correspond directly to the server module number. For example, server module 2 is connected to port 2 on the Gigabit Ethernet pass-through module. To create a private network between server module 2 and server module 4 (see Figure 1-1) using a Gigabit pass-through module, connect a standard Ethernet cable to port 2 and port 4 on the Gigabit Ethernet pass-through module.

Table 2-2 provides a cable connection matrix for configuring the private network using a Dell PowerConnect™ Ethernet switch or Gigabit Ethernet pass-through modules in your PowerEdge server enclosure.

Table 2-2. Cabling the Private Network

Module Type	Cluster Configuration	Private Network Configuration
Ethernet switch module	Two or more nodes in one PowerEdge server enclosure.	The private network is established using internal connections in the system chassis.
	Two or more nodes in two PowerEdge server enclosures.	<ol style="list-style-type: none">1 Identify the Ethernet switch module on each PowerEdge server enclosure that is connected to the private network-configured adapters on each server module.2 Connect a standard Ethernet cable from the appropriate switch on system 1 to the appropriate switch on system 2.

Table 2-2. Cabling the Private Network (continued)

Module Type	Cluster Configuration	Private Network Configuration
Gigabit Ethernet pass-through module	Two nodes in one or two PowerEdge server enclosure(s).	Connect a standard Ethernet cable to the corresponding cluster node ports on the Ethernet pass-through module. OR Connect a standard Ethernet cable from the corresponding cluster node ports on each Ethernet pass-through module to an external switch.
	Three or more nodes in one or two PowerEdge server enclosure(s).	Connect a standard Ethernet cable from the corresponding cluster node ports on the Ethernet pass-through module to an external switch.

Cabling the Public Network

The public (client) network is used for client access and private network failover. Each cluster node connects to the public network through the first embedded NIC port on the server module system board. The first NIC port connects to the internal connections in the system chassis to the first Ethernet switch module or the associated port in the corresponding Gigabit pass-through module. The Ethernet switch or Gigabit pass-through module is connected to an external switch in the public network to provide client access to the server module.



NOTE: It is recommended that you configure the first NIC port on each server module for the public network. If additional NIC ports are available on your server module, you can use NIC teaming to aggregate multiple ports for the public network.

Cabling the Storage Systems

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



NOTE: The cluster nodes represent the individual server modules in the PowerEdge server enclosure. The designations namely, node 1 and node 2 and the server enclosure are used for illustrative purposes only. The system and the storage components in your cluster configuration may vary.

Direct-Attached Cluster

In a direct-attached cluster configuration, the redundant Fibre Channel/iSCSI HBA ports are connected to a supported storage system using pass-through modules. Two cluster nodes in a direct-attached configuration can be configured within one PowerEdge server enclosure or configured between two PowerEdge server enclosures. Direct-attached configurations are self-contained and do not share any physical resources with other server or storage systems outside of the cluster.

Each server module is labeled with an identification (ID) number located above each module on the front of the PowerEdge server enclosure that corresponds with the identical port number on the pass-through module. For example, server module 4 connects to port 4 on the pass-through module.

To connect your server modules to a storage system in a direct-attached configuration:

- 1 Identify the server modules that you wish to configure in a direct-attached configuration.
- 2 Locate the identification number for each server module you selected in step 1.
- 3 Match the server module identification numbers with the corresponding port numbers on the pass-through module.
- 4 Connect cables from the selected ports to the appropriate SPs in the supported storage system.



NOTE: Figure 2-2 illustrates a PowerEdge 1855/1955 server enclosure configured with a Dell|EMC CX3-10c storage enclosure. You can use a similar configuration model for other supported server modules, server enclosures, and storage arrays.

Figure 2-2 shows an example of a direct-attached, two-node cluster configuration using redundant connections to server modules 5 and 6.

See your Dell PowerEdge system documentation for more information about the Fibre Channel/Ethernet pass-through module.

Cabling One PowerEdge Cluster to a Dell|EMC CX3-10c Storage System


In this configuration, each server module attaches to the storage system using CAT5e or CAT6 LAN cables with RJ45 connectors that attach to Gigabit Ethernet daughter cards in the server modules and the Gigabit iSCSI SP ports in the Dell|EMC storage system.

To connect two server modules (for example, module 4 and 5) to a storage system in a direct-attached configuration:

- 1** Identify the server modules that you wish to configure in a direct-attached configuration.
- 2** Locate the identification number for each server module selected in step 1.
- 3** Match the server module identification numbers with the corresponding port numbers on the Ethernet pass-through module.
- 4** Connect server module 4 to the storage system.
 - a** Connect a LAN cable from port-4 of pass-through module 1 to SP-A port 1 iSCSI.
 - b** Connect a LAN cable from port-4 of pass-through module 2 to SP-B port 0 iSCSI.
- 5** Connect server module 5 to the storage system.
 - a** Connect a LAN cable from port-5 of pass-through module 1 to SP-B port 1 iSCSI.
 - b** Connect a LAN cable from port-5 of pass-through module 2 to SP-A port 0 iSCSI.

Figure 2-2 shows an example of a direct-attached, two-node iSCSI cluster configuration using redundant connections to server modules 5 and 6 in PowerEdge 1855/1955 server enclosure connected to Dell|EMC iSCSI storage system.

Figure 2-2. Direct-Attached Cluster Configuration Using CX3-10c Storage System



- | | | | |
|---|--------------------------------|----|--------------------------------|
| 1 | Ethernet pass-through module 2 | 7 | iSCSI ports |
| 2 | port 5 | 8 | port 4 |
| 3 | port 4 | 9 | port 5 |
| 4 | SP-B | 10 | Ethernet pass-through module 1 |
| 5 | SP-A | 11 | internal connections |
| 6 | Fibre Channel ports | | |

Network-Attached Cluster

A network-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. Network-attached cluster configurations provide flexibility, expandability, and performance with either iSCSI or Fibre Channel storage arrays.

Network-Attached Cluster Configured With an Internal Fabric

An internal fabric incorporates switch modules that do not connect to external switches. Each server module I/O port is internally connected to a port on a switch module. Up to two storage systems can be connected to the outgoing ports on a switch module with internal fabric.

Cablings One PowerEdge Server Enclosure to One Storage System


In this configuration, a PowerEdge M1000e server enclosure is connected to Dell|EMC CX3-20 storage system. The server and storage components in your configuration may vary.

To cable one PowerEdge server enclosure to Dell|EMC CX3-20 storage system:

- 1 Connect a cable from Fibre Channel switch module 1 to SP-A fibre port 0 (first fibre port).
- 2 Connect a cable from Fibre Channel switch module 1 to SP-B fibre port 1 (second fibre port).
- 3 Connect a cable from Fibre Channel switch module 2 to SP-A fibre port 1 (second fibre port).
- 4 Connect a cable from Fibre Channel switch module 2 to SP-B fibre port 0 (first fibre port).

Figure 2-3 shows how to cable a PowerEdge M1000e sever enclosure to the CX3-20 storage system. The server module's dual-port Fibre Channel daughter card and Fibre Channel switch modules are connected internally in the system chassis.

Figure 2-3. Cabling One PowerEdge M1000e Server Enclosure to the CX3-20 Storage Processor Enclosure (SPE)



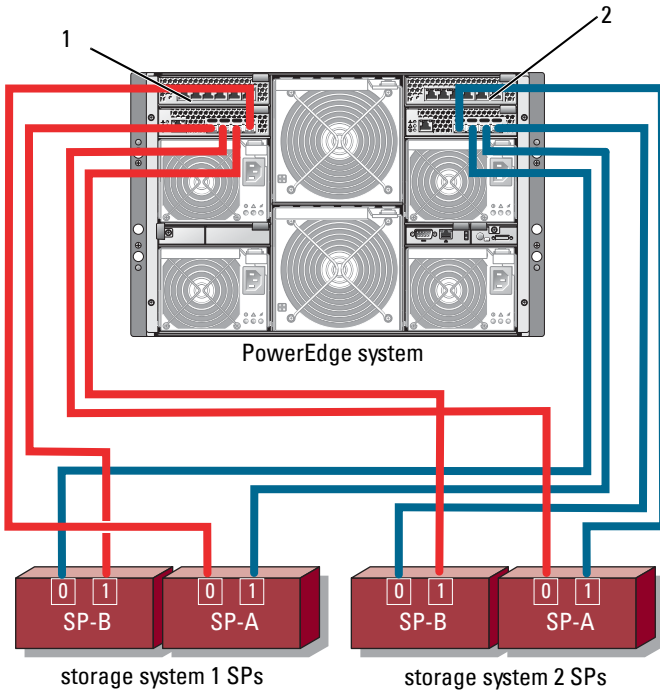
- | | | | |
|---|---|---|---|
| 1 | Fibre Channel switch module in I/O bay C1 | 5 | Fibre Channel switch module in I/O bay C2 |
| 2 | SP-B | 6 | SP-A |
| 3 | 0Fibre | 7 | 0Fibre |
| 4 | 1Fibre | 8 | 1Fibre |

Cabling One PowerEdge Server Enclosure to Multiple Storage Systems

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

In Figure 2-4, a PowerEdge 1855/1955 server enclosure is connected to the Dell|EMC iSCSI storage systems. The server and storage components in your configuration may vary.

Figure 2-4. One PowerEdge 1855/1955 Server Enclosure Cabled to Two Storage Systems



- 1 Gigabit Ethernet switch module 1 2 Gigabit Ethernet switch module 2

NOTE: If you are using a CX3-20c or CX3-40c storage system for the configuration in Figure 2-4, you can also connect the remaining iSCSI ports 2 *iSCSI* and 3 *iSCSI* (not shown) depending on the required level of redundancy.

Cabling One PowerEdge Server Enclosure to a Tape Library

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

Figure 2-5 shows a supported PowerEdge cluster configuration using redundant Fibre Channel switches and a tape library. In this configuration, each of the cluster nodes can access the tape library to provide backup for your local disk resources, as well as your cluster disk resources.


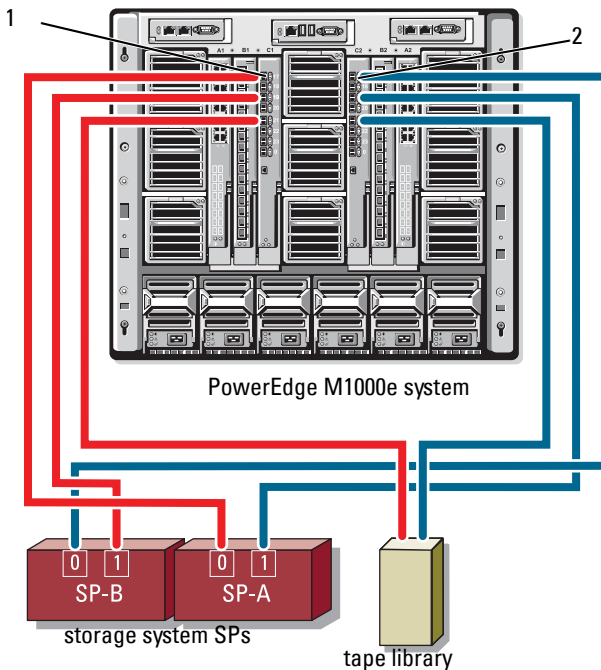
 **NOTE:** While tape libraries can be connected to multiple fabrics, they do not provide path failover.

Figure 2-5. Cabling a Storage System and a Tape Library



- 1 Fibre Channel switch module in I/O bay C1
- 2 Fibre Channel switch module in I/O bay C2

For more information on configuring the tape and storage components, see the storage and tape backup documentation.

Cabling Two PowerEdge Server Enclosures to One or Two Dell|EMC Storage Systems

The PowerEdge blade cluster supports up to two PowerEdge server enclosures and up to two Dell|EMC storage systems. In this configuration, two inter-switch links (ISLs) are required to connect each pair of Fibre Channel switch modules between the PowerEdge server enclosures.

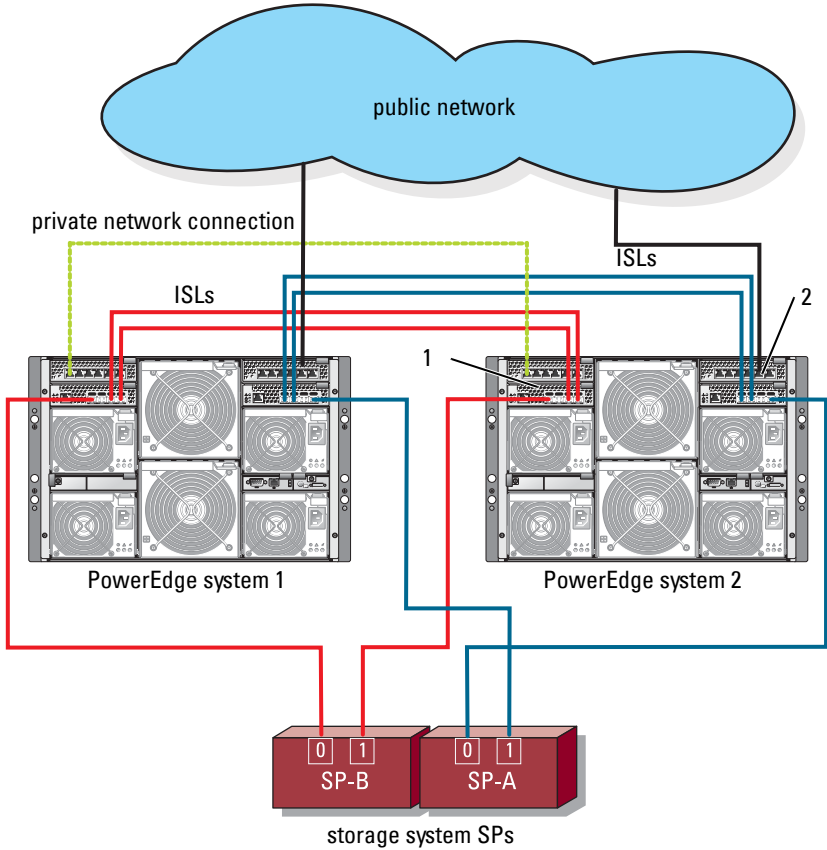
For high availability, each PowerEdge server enclosure requires direct paths to the attached storage systems. This configuration ensures that all running applications are available if one of the PowerEdge server enclosures needs to be shut down for maintenance.

Figure 2-6 shows two PowerEdge 1855/1955 server enclosures cabled to one supported Dell|EMC storage system in an iSCSI configuration.

Figure 2-7 shows two PowerEdge M1000e server enclosures cabled to two supported Dell|EMC storage systems in Fibre Channel configuration.

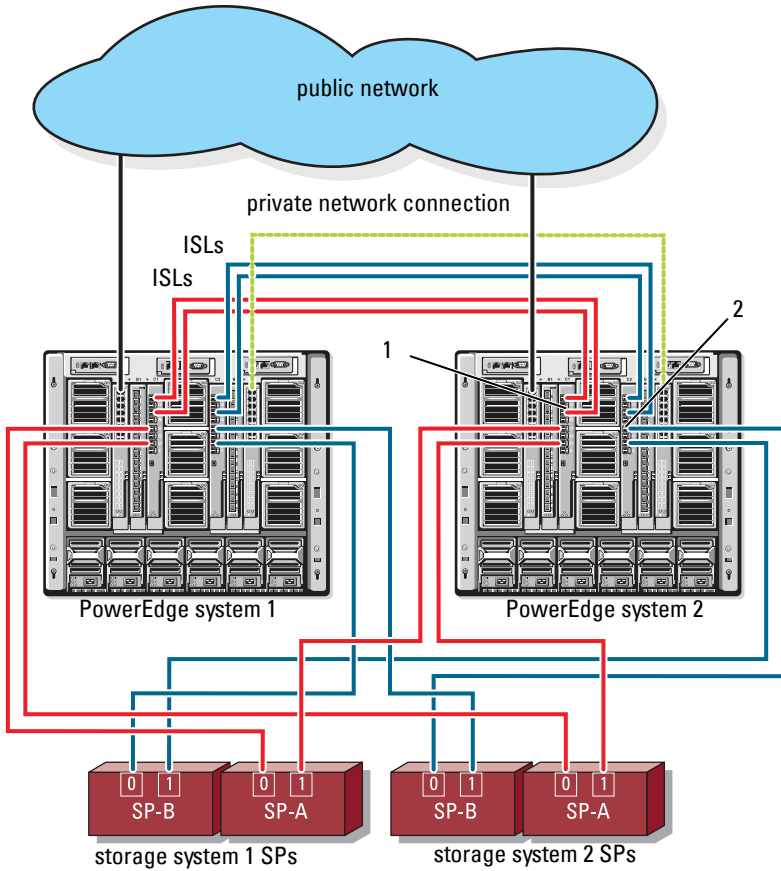
In both figures, the server module's dual-port Fibre Channel daughter card and Fibre Channel modules are connected internally in the system chassis. You can use a similar configuration model for other supported server modules, server enclosures, and storage arrays.

Figure 2-6. Cluster Configuration Using Two PowerEdge Server Enclosures and One Storage System



- 1 Ethernet switch module for iSCSI traffic (2)
- 2 Ethernet switch module for cluster network (2)

Figure 2-7. Two PowerEdge Server Enclosures Connected to Two Storage Systems



- 1 Fibre Channel switch module (2)
- 2 Fibre Channel switch module (2)

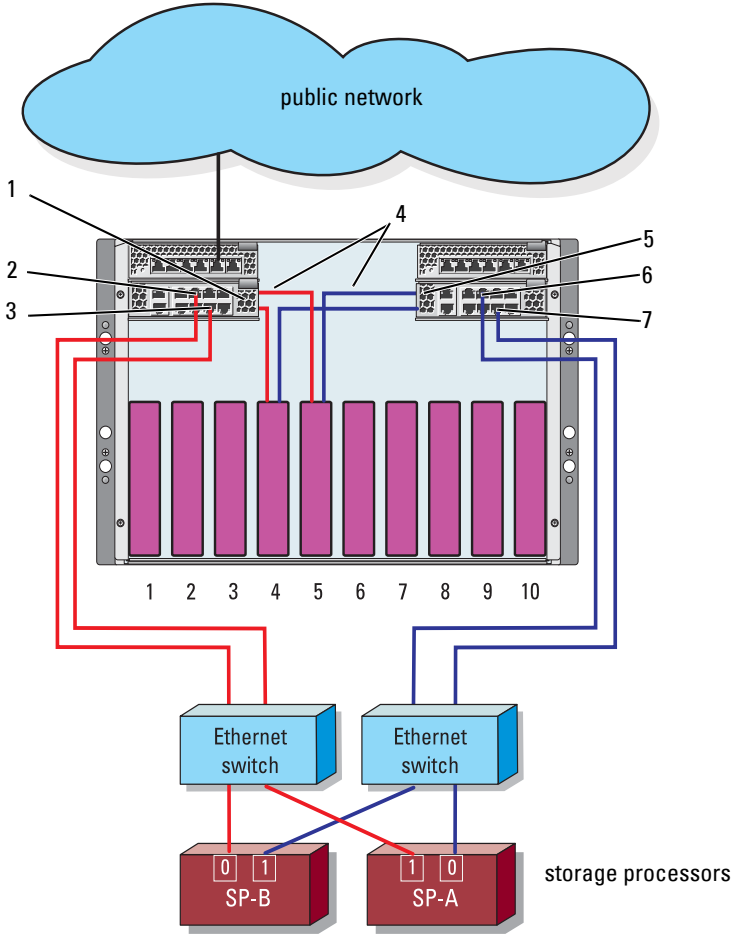
Network-Attached Cluster Connected to an External Fabric

External Switches incorporate the external switch network or SAN connected to the PowerEdge Enclosure through the pass-through modules or switch modules. These configurations allow you to extend the Fibre Channel/iSCSI network by connecting more storage systems. The following sections provide examples for these configurations.


Cabling PowerEdge Server Enclosures With Pass-Through Modules to an External Fabric

Figure 2-8 shows an example of a switch-attached iSCSI cluster with embedded Ethernet pass-through modules that are connected to external Ethernet switches (dedicated for iSCSI). Figure 2-8 illustrates a PowerEdge 1855/1955 server enclosure configured with an iSCSI storage array. You can use a similar configuration model for other supported server modules, server enclosures, and storage arrays.

Figure 2-8. External Network-Attached iSCSI Cluster With Embedded Ethernet Pass-Through Modules



- | | | | |
|---|------------------------------|---|------------------------------|
| 1 | Ethernet pass-through module | 5 | Ethernet pass-through module |
| 2 | port 5 | 6 | port 5 |
| 3 | port 4 | 7 | port 4 |
| 4 | Internal connections | | |

 **NOTE:** Figure 2-8 is for illustration purposes only. Some cluster connections are routed internally through the PowerEdge server enclosure. The server and storage components in your cluster configuration may vary.

Cabling PowerEdge Server Enclosures With Embedded Switch Modules to an External Fabric

Figure 2-9 shows an example of a PowerEdge M1000e system cluster with embedded Fibre Channel switch modules connected to an external SAN-attached fabric.


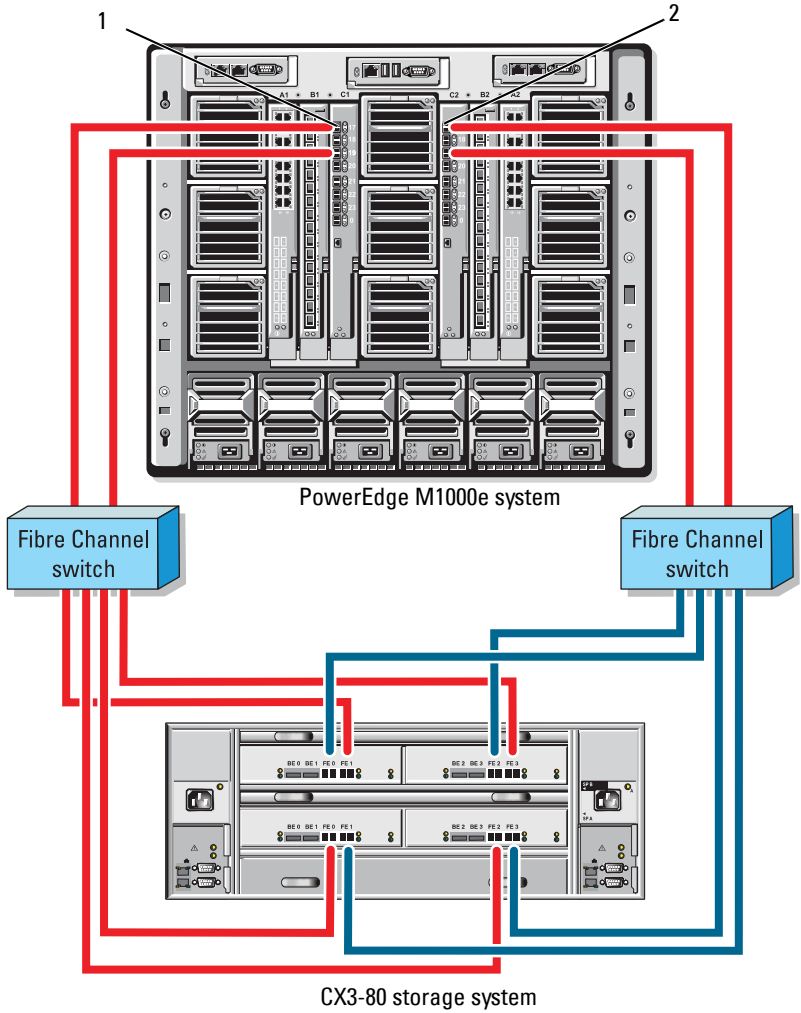
 **NOTE:** Figure 2-9 is for illustration only. Some cluster connections are routed internally through the PowerEdge server enclosure. The server and storage components in your cluster configuration may vary.

Figure 2-9. External Network-Attached Cluster With Embedded Fibre Channel Switch Modules



- 1 Fibre Channel switch module in I/O bay C1
- 2 Fibre Channel switch module in I/O bay C2

Installing the Operating System and Configuring Your Cluster

For information about installing the operating system, Microsoft Cluster Services, and configuring your cluster, see the *Installation and Troubleshooting Guide* for your cluster system located on the Dell Support website at support.dell.com.

Maintaining Your Cluster

For information about maintaining your cluster, see the *Installation and Troubleshooting Guide* for your cluster system located on the Dell Support website at support.dell.com.

