

Dell™ PowerEdge™ Cluster  
FE500W Systems  
**Platform Guide**



# Notes, Notices, and Cautions



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



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



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

---

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

© 2005–2007 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *OpenManage*, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc.; *Microsoft*, *Windows*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries; *EMC*, *Navisphere*, and *PowerPath* are registered trademarks and *MirrorView* is a trademark of EMC Corporation.

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

# Contents

<b>Supported Cluster Configurations</b> . . . . .	<b>5</b>
HBA Support for PowerEdge Cluster FE500W Configurations . . . . .	7
Fibre Channel Switches . . . . .	9
Supported Combinations of Switch/HBA. . . . .	9
Rules and Guidelines . . . . .	10
<b>Installing Peripheral Components in Your Cluster Node</b> . . . . .	<b>10</b>
<b>Attaching Your Cluster to a Shared Storage System in a Direct-Attached Cluster Configuration</b> . . . . .	<b>20</b>
Rules and Guidelines . . . . .	20
<b>Attaching Your Cluster to a Shared Storage System in a SAN Configuration</b> . . . . .	<b>22</b>
Rules and Guidelines . . . . .	22
<b>Dell Cluster Configuration Support Matrix</b> . . . . .	<b>24</b>
<b>Best Practices</b> . . . . .	<b>25</b>
Using a Tape Backup Library in a SAN. . . . .	25
The Cluster Disks are not Initialized in Disk Management . . . . .	26





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

This document provides information for installing and connecting peripheral hardware, storage, and storage area network (SAN) components for your Dell™ PowerEdge™ Cluster FE500W solution. The configuration information in this document is specific to the Microsoft® Windows® 2000 Advanced Server and Microsoft Windows Server® 2003 operating systems.

This document covers the following topics:

- Supported cluster configurations
- Installing peripheral components in your cluster node
- Attaching your cluster to a shared storage system in a direct-attached cluster configuration
- Attaching your cluster to a shared storage system in a SAN configuration
- Dell cluster configuration support matrix
- Best practices



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



**NOTE:** Configurations not listed in this document may not be certified or supported by Dell or Microsoft. For more information on component revisions that are certified, see the Dell High Availability Clustering website at [www.dell.com/ha](http://www.dell.com/ha).


## Supported Cluster Configurations


This section provides information about supported cluster configurations for your PowerEdge cluster solution.


Table 1-1 provides a list of supported cluster configurations for the PowerEdge Cluster FE500W systems running Windows 2000 Advanced Server or Windows Server 2003.


**Table 1-1. Supported Cluster Configurations**


<b>PowerEdge Cluster</b>	<b>Supported PowerEdge Systems</b>	<b>Supported Storage Systems</b>	<b>Supported Cluster Interconnect (for the Private Network)</b>
FE500W	1650, 1750, 1800, 1850, 1855, 1950, 1955, 2500, 2550, 2600, 2650, 2800, 2850, 2900, 2950, 2970, 4400, 4600, 6400, 6450, 6600, 6650, 6800, 6850, 6950, 8450, R900, M600, and M605.	Dell   EMC CX700 Dell   EMC CX500 Dell   EMC CX300	Any network interface card (NIC) supported by the system.  <b>NOTE:</b> All nodes in a cluster must use identical NICs for the cluster interconnect.  <b>NOTE:</b> NIC teaming is only supported for the public (client) network, and should not be configured on the private network.

 **NOTE:** Windows Server 2003 Enterprise x64 Edition and Windows Server 2003 R2 Enterprise x64 Edition do not support a rolling upgrade from Windows 2000 Advanced Server.

 **NOTE:** All cluster nodes in a PowerEdge Cluster FE500W solution must be of the same model and must run the same operating system. Mixing Windows 2000 Advanced Server and Windows Server 2003 in the same cluster is not supported except during a rolling upgrade.

 **NOTE:** Reference to PowerEdge 1950, 2900, and 2950 in this document also implies reference to PowerEdge 1950 III, 2900 III, and 2950 III respectively.

 **NOTE:** Windows Server 2003 Enterprise x64 Edition and Windows Server 2003 R2 Enterprise x64 Edition are only supported on the PowerEdge 1800, 1850, 1855, 1950, 1955, 2800, 2850, 2900, 2950, 2970, 6800, 6850, 6950, R900, M600, and M605 systems.

 **NOTE:** Windows 2000 Advanced Server is not supported on the PowerEdge 1950, 2900, 2950, 2970, 6950, R900, M600, and M605 systems.

### **Obtaining More Information**

For a detailed list of related documentation, see the *Dell PowerEdge Cluster FE500W Systems Installation and Troubleshooting Guide*. You can also see the Dell Support website at [support.dell.com](http://support.dell.com).

## HBA Support for PowerEdge Cluster FE500W Configurations

Table 1-2 lists the systems and the HBAs that are supported for PowerEdge Cluster FE500W configurations running Windows 2000 Advanced Server or Windows Server 2003.

Two single-port HBAs must be installed in each node, except when using the embedded dual-port HBA for PowerEdge 6850 or when using the QLogic QLE2462 dual-port HBA in supported Dell PowerEdge systems. Each node in the cluster must have homogenous (identical) HBAs.

For a list of recommended installation locations for embedded and peripheral HBAs, see "Installing Peripheral Components in Your Cluster Node" on page 10.



**NOTE:** The HBAs installed in clusters using redundant paths must be identical. Cluster configurations are tested and certified using identical HBAs installed in all of the cluster nodes. Using dissimilar HBAs in your cluster nodes is not supported.

**Table 1-2. Supported HBAs for PowerEdge Cluster FE500W Configurations**

PowerEdge System	Emulex LP982 or LP9802 (PCI-X) HBA	QLogic QLA2340 (PCI-X) HBA	Emulex LP1050 EX (PCI Express [PCIe]) HBA	Emulex LPe1150 (PCIe) HBA	QLogic QLE2360 (PCIe) HBA	QLogic QLE2460 (PCIe) HBA	QLogic QLE246 2 dual-port (PCIe) HBA	Emulex LP10000 (PCI-X) HBA	Dell 2362M adapter
1650	X	X							
1750	X	X						X	
1800	X	X	X	X	X	X	X	X	
1850	X*	X*	X**	X**	X**	X**	X**	X*	
1950		X*		X**		X**	X**	X*	
2500	X	X							
2550		X							
2600/2650	X	X						X	
2800	X	X	X	X	X	X	X	X	
2850	X*	X*	X**	X**	X**	X**	X**	X*	
2900		X		X		X	X	X	
2950		X*		X**		X**	X	X*	

**Table 1-2. Supported HBAs for PowerEdge Cluster FE500W Configurations (continued)**

PowerEdge System	Emulex LP982 or LP9802 (PCI-X) HBA	QLogic QLA2340 (PCI-X) HBA	Emulex LP1050 EX (PCI Express [PCIe]) HBA	Emulex LPe1150 (PCIe) HBA	QLogic QLE2360 (PCIe) HBA	QLogic QLE2460 (PCIe) HBA	QLogic QLE246 2 dual-port (PCIe) HBA	Emulex LP10000 (PCI-X) HBA	Dell 2362M adapter
2970				X		X	X		
4400	X	X							
4600	X	X							X
6400/6450	X	X							
6600/6650	X	X							X
6800	X	X	X	X	X	X	X	X	
6850	X	X	X	X	X	X	X	X	X
6950				X		X	X		
8450	X	X							
R900				X		X	X		

\* The PowerEdge system must have a PCI-X riser installed in order to use this HBA.

\*\* The PowerEdge system must have a PCIe riser installed in order to use this HBA.

**Table 1-3. Supported HBAs for PowerEdge Blade Cluster Configurations**

PowerEdge System	Dell 2362M adapter	Qlogic QME2462	Emulex LPe 1105M*	QLogic QME2472	Emulex LPe11002
1855	X	X	X		
1955	X	X	X		
M600				X	X
M605				X	X

\* Only supported with Dell|EMC CX series storage.



## Fibre Channel Switches

- Dual (redundant) fabric configurations are required.
- A maximum of 16 switches may be used in a SAN.
- A minimum of two and a maximum of eight Inter-Switch Links (ISLs) may exist between any two directly communicating switches. A single ISL is permitted only when connecting to a remote switch in an EMC<sup>®</sup> MirrorView™ configuration.
- A maximum of three hops (the number of ISLs each data frame must traverse) may exist between a host and a storage system.

**Table 1-4. Supported Fibre Channel I/O Modules for PowerEdge Blade Servers**

PowerEdge System	Brocade 3014 Brocade 4016	McData 4314 McData 4416	Dell Fibre Channel Pass-Through	Brocade SW4424	Emulex PT1016 Pass-Through
1855	X	X	X		
1955	X	X	X		
M600				X	X
M605				X	X

## Supported Combinations of Switch/HBA

Table 1-5 lists the combinations of the switches and the HBAs that are supported for PowerEdge Cluster FE500W configurations running Windows 2000 Advanced Server or Windows Server 2003.

**Table 1-5. Supported Combinations of Switch/HBA**

Switch/HBA	LP982/ LP9802	Dell 2362M/LP1000/ LP1050-EX/QLA2340/ QLE2360	LPe1150/ QLE2460/ QLE2462
Brocade SilkWorm 3200/3800	X	X	
Brocade SilkWorm 3250/3850/4100 McData Sphereon 4500	X	X	X
Brocade SilkWorm 200E McData Sphereon 4400/4700		X	X

## Rules and Guidelines

When configuring your cluster, all cluster nodes must contain identical versions of the following:

- Operating systems and service packs
- Hardware, drivers, firmware, or BIOS for the NICs, HBAs, and any other embedded or peripheral hardware components
- Dell OpenManage™ Server Administrator systems management software and EMC Navisphere® storage management software

### Maximum Distance Between Cluster Nodes

The maximum cable length allowed from an HBA to a switch, an HBA directly connected to a storage system, or a switch to a storage system is 300 meters using multimode fibre at 2 Gbps or 100 meters using multimode fibre at 4 Gbps. The total distance between an HBA and a storage system may be increased through the use of switch ISLs.

The maximum cable length for Fast Ethernet and copper Gigabit Ethernet is 100 meters, and for optical Gigabit Ethernet is 550 meters. This distance may be extended using switches and the virtual local area network technology. The maximum latency for a round-trip network packet between nodes is 500 milliseconds.

### Obtaining More Information

For installation instructions for the hardware configurations running Windows 2000 Advanced Server or Windows Server 2003, see the *Dell PowerEdge Cluster FE500W Systems Installation and Troubleshooting Guide*.

## Installing Peripheral Components in Your Cluster Node

In this section, Table 1-6 provides information about the PCI slot configurations and Table 1-7 provides information about the PCI slot assignments.



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

**Table 1-6. PCI Expansion Slot Configurations for PowerEdge Cluster Nodes**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>Slot</b>	<b>Slot Type</b>	<b>Slot Speed</b>
1650	Any	1	PCI	64-bit, 66 MHz or 32-bit, 33 MHz
		2	PCI	64-bit, 66 MHz
1750	Any	1	PCI-X or PCI	64-bit, 133 MHz PCI-X or 64-bit, 33 MHz PCI
		2	PCI-X	64-bit, 133 MHz
1800	N/A	1	PCI	64-bit, 66 MHz
		2	PCIe	2.5 GHz x4-lane width
		3	PCIe	2.5 GHz x8-lane width
		4	PCI	32-bit, 33 MHz
		5-6	PCI-X	64-bit, 100 MHz
1850	Standard	1	PCI-X	64-bit, 133 MHz
		2	PCI-X	64-bit, 100 MHz
	PCI-X with RAID on motherboard (ROMB)	1	PCI-X	64-bit, 133 MHz
		2	PCI-X	64-bit, 100 MHz
	PCIe with ROMB	1	PCIe	2.5 GHz x4-lane width
		2	PCIe	2.5 GHz x8-lane width
1950	PCI-X	1-2	PCI-X	64-bit, 133 MHz
	PCIe	1-2	PCIe	2.5 GHz x8-lane width
2500	N/A	1-2	PCI	64-bit, 66 MHz
		3-5	PCI	64-bit, 33 MHz
		6-7	PCI	32-bit, 33 MHz

**Table 1-6. PCI Expansion Slot Configurations for PowerEdge Cluster Nodes (continued)**

PowerEdge System	Riser Board Option	Slot	Slot Type	Slot Speed		
2550	N/A	1-3	PCI	64-bit, 33 MHz		
2600	N/A	1	PCI	32-bit, 33 MHz		
		2-5	PCI-X	64-bit, 100 MHz		
		6-7	PCI-X	64-bit, 133 MHz		
2650	N/A	1	PCI-X	64-bit, 100 MHz		
		2-3	PCI-X	64-bit, 133 MHz		
<b>NOTE:</b> Slot 1 must be empty for slot 2 to attain an operating speed of 133 MHz.						
2800	N/A	1	PCI	32-bit, 33 MHz		
		2-5	PCI-X	64-bit, 133 MHz		
		6	PCIe	2.5 GHz x4-lane width		
		7	PCIe	2.5 GHz x8-lane width		
2850	PCI-X	1-3	PCI-X	64-bit, 133 MHz		
		<b>NOTE:</b> If slot 1 is populated, slots 2 and 3 operate at 100 MHz.				
		PCIe	1	PCIe	2.5 GHz PCIe x4-lane width	
		2	PCIe	2.5 GHz PCIe x8-lane width		
		3	PCI-X	64-bit, 100 MHz		
2900	N/A	1-2	PCI-X	64-bit, 133 MHz		
		3	PCIe	2.5 GHz x8-lane width		
		4-6	PCIe	2.5 GHz x4-lane width		
2950	PCI-X	1	PCIe	2.5 GHz x8-lane width		
		2-3	PCI-X	64-bit, 133MHz		
	PCIe	1-2	PCIe	2.5 GHz x8-lane width		
		3	PCIe	2.5 GHz x4-lane width		

**Table 1-6. PCI Expansion Slot Configurations for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>Slot</b>	<b>Slot Type</b>	<b>Slot Speed</b>
2970	N/A	1	PCIe	2.5 GHz x4-lane width
		2-3	PCIe	2.5 GHz x8-lane width
4400	N/A	1-2	PCI	64-bit, 66 MHz
		3-6	PCI	64-bit, 33 MHz
		7	PCI	32-bit, 33 MHz
4600	N/A	1	PCI	32-bit, 33 MHz
		2-3	PCI-X	64-bit, 100 MHz
		4-5	PCI-X	64-bit, 100 MHz
		6-7	PCI-X	64-bit, 100 MHz
6400	N/A	1	PCI	32-bit, 33 MHz
		2-5	PCI	64-bit, 33 MHz
		6-7	PCI	64-bit, 66 MHz
6450	N/A	1	PCI	32-bit, 33 MHz
		2-5	PCI	64-bit, 33 MHz
		6-7	PCI	64-bit, 66 MHz
6600	N/A	1	PCI	32-bit, 33 MHz
		2-3	PCI-X	64-bit, 100 MHz
		4-5	PCI-X	64-bit, 100 MHz
		6-7	PCI-X	64-bit, 100 MHz
		8-9	PCI-X	64-bit, 100 MHz
		10-11	PCI-X	64-bit, 100 MHz

**Table 1-6. PCI Expansion Slot Configurations for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>Slot</b>	<b>Slot Type</b>	<b>Slot Speed</b>
6650	N/A	1	PCI	32-bit, 33 MHz
		2-3	PCI-X	64-bit, 100 MHz
		4-5	PCI-X	64-bit, 100 MHz
		6	PCI-X	64-bit, 100 MHz
		7	PCI-X	64-bit, 100 MHz
		8	PCI-X	64-bit, 100 MHz
6800	N/A	1	PCIe	2.5-GHz x8-lane width
		2	PCI-X	64-bit, 133 Mhz
		3-5	PCIe	2.5-GHz x4-lane width
		6-7	PCI-X	64-bit, 100 Mhz
6850	N/A	1	PCIe	2.5-GHz x8-lane width
		2	PCI-X	64-bit, 133 Mhz
		3-5	PCIe	2.5-GHz x4-lane width
		6-7	PCI-X	64-bit, 100 Mhz
6950	N/A	1-2	PCIe	2.5-GHz x4-lane width
		3-4	PCIe	2.5-GHz x8-lane width
		5-7	PCIe	2.5-GHz x4-lane width
8450	N/A	1-2	PCI	64-bit, 33 MHz
		3-6	PCI	64-bit, 33 MHz
		7-8	PCI	64-bit, 66 MHz
		9-10	PCI	64-bit, 66 MHz
R900	N/A	1-4	PCIe	2.5-GHz x8-lane width
		5-7	PCIe	2.5-GHz x4-lane width

**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
1650	Any	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
1750	Any	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI or PCI-X slots.
1800	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI/PCI-X slots 1, 5, or 6.
		Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBAs in PCIe slots 2 and 3.
1850	Standard	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 1 and 2.
	PCI-X with raid on motherboard (ROMB)	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 1 and 2.

**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
	PCIe with ROMB	Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1 and/or 2.
1950	PCI-X	Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 1 and 2.
	PCIe	Emulex LP1150 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1 and/or 2.
2500	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBAs in PCI slots 1, 2, 3, 4, or 5.
2550	N/A	QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
2600	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in PCI-X slots 2, 3, 4, or 5.
2650	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI or PCI-X slots.



**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
2800	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 2, 3, 4, or 5.
		Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 6 and/or 7.
2850	PCI-X	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI-X slots.
	PCIe	Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1 and/or 2.
2900	N/A	Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 1 and 2.
		Emulex LPe1150 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in any of the available PCIe slots.

**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
2950	PCI-X	Emulex LP10000 QLogic QLA2340	Install the HBAs in PCI-X slots 2 and 3.
	PCIe	Emulex LPe1150 QLogic QLE2460	Install the HBA(s) in any of the available PCIe slots.
	Any	QLogic QLE2462	Install the HBA in any available PCIe slots.
2970	N/A	Emulex LPe1150 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in any of the available PCIe slots.
4400	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
4600	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI or PCI-X slots.
6400	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
6450	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
6600	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI or PCI-X slots.

**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
6650	N/A	Emulex LP982 Emulex LP9802 Emulex LP10000 QLogic QLA2340	Install the HBA(s) in any of the available PCI or PCI-X slots.
6800	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in PCI-X slots 2, 6, or 7.
		Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1, 3, 4, or 5.
6850	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in PCI-X slots 2, 6, or 7.
		Dell 2362M	For installation instructions, see the <i>Dell PowerEdge 6850 Systems Installation and Troubleshooting Guide</i> .
		Emulex LP1050-EX Emulex LPe1150 QLogic QLE2360 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1, 3, 4, or 5.
6950	N/A	Emulex LPe1150 QLogic QLE2460 QLogic QLE2462	Install HBA(s) in PCIe slots 1, 2, 3, 4, 5, 6, or 7.

**Table 1-7. PCI Expansion Slot Assignments for PowerEdge Cluster Nodes (continued)**

<b>PowerEdge System</b>	<b>Riser Board Option</b>	<b>HBA</b>	<b>PCI Slot Assignment</b>
8450	N/A	Emulex LP982 Emulex LP9802 QLogic QLA2340	Install the HBA(s) in any of the available PCI slots.
R900	PCIe	Emulex LPe1150 QLogic QLE2460 QLogic QLE2462	Install the HBA(s) in PCIe slots 1, 2, 3, or 4.

**NOTE:** Whenever possible, it is recommended that the HBAs be placed on separate buses to balance the load on the system. These buses are identified as separate rows in Table 1-6.

## Attaching Your Cluster to a Shared Storage System in a Direct-Attached Cluster Configuration

This section provides the rules and guidelines for attaching your cluster nodes to the shared storage system using a direct connection (without Fibre Channel switches).

In a direct-attached configuration, both cluster nodes are connected directly to the storage system.

### Rules and Guidelines

The rules and guidelines described in Table 1-8 apply to direct-attached clusters.

**Table 1-8. Direct-Attached Clusters Rules and Guidelines**

<b>Rule/Guideline</b>	<b>Description</b>
Operating system	Each direct-attached cluster must run Windows 2000 Advanced Server or Windows Server 2003.
Windows 2000 Advanced Server service pack	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Windows Server 2003 service pack	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Primary storage	Each Windows 2000 Advanced Server and Windows Server 2003 cluster can support up to 22 unique drive letters for shared logical drives. Windows Server 2003 can support additional physical drives through mount points. <b>NOTE:</b> You can attach only one storage system directly to the cluster.
Fibre Channel HBAs supported	To determine which HBAs are supported on your PowerEdge system, see Table 1-2.
Fibre Channel HBA driver versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Fibre Channel HBA firmware versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Dell EMC CX300, CX500, or CX700 core software	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
EMC PowerPath <sup>®</sup> versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.

# Attaching Your Cluster to a Shared Storage System in a SAN Configuration

This section provides the rules and guidelines for attaching your PowerEdge cluster nodes to the shared storage systems through a Dell|EMC SAN using redundant Fibre Channel switch fabrics.

For a cluster based on PowerEdge 1855, 1955, M600, or M605 the following hardware components in a SAN configuration may be used:

- Multiple clusters and non-clustered server modules, but is limited to no more than two PowerEdge 1855/1955 server enclosures.
- The supported storage system configurations using internal fabrics include:
  - One Dell|EMC storage system
  - Two Dell|EMC storage systems
  - One Dell|EMC storage system and one Dell PowerVault™ tape library
- With external SAN, a system can support up to four attached storage systems and up to two tape libraries.

## Rules and Guidelines

The rules and guidelines described in Table 1-9 apply to SAN-attached clusters.

**Table 1-9. SAN-Attached Clusters Rules and Guidelines**

Rule/Guideline	Description
Operating system	Each cluster attached to the SAN must run Windows 2000 Advanced Server or Windows Server 2003.
Windows 2000 Advanced Server service pack	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Windows Server 2003 service pack	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.

**Table 1-9. SAN-Attached Clusters Rules and Guidelines (continued)**

<b>Rule/Guideline</b>	<b>Description</b>
Primary storage	<p>Each Windows 2000 Advanced Server and Windows Server 2003 cluster can support up to 22 unique drive letters for shared logical drives. Windows Server 2003 can support additional physical drives through mount points.</p> <p>Up to four Dell EMC Fibre Channel disk arrays are supported per cluster in a SAN environment.</p> <p>Up to two Dell EMC Fibre Channel disk arrays are supported in a configuration using only embedded switches.*</p>
Secondary storage	<p>Up to two Dell PowerVault™ tape libraries that are supported by your cluster nodes in a Dell EMC SAN environment may be used by the cluster.</p> <p>Any system attached to the SAN can share these devices.</p>
Fibre Channel switch configuration	Redundant switch fabrics required
Fibre Channel switch zoning	Single-initiator zoning
Fibre Channel module	<p>Two Brocade SilkWorm 3014/4016 Fibre Channel switch modules.</p> <p>Two McData 4314/4416 Fibre Channel switch modules.</p> <p>Two Fibre Channel pass-through modules.</p>
Fibre Channel switch (external SAN)	<p>EMC DS-220B-DE (Brocade SilkWorm 200E)</p> <p>EMC DS-8B2 (Brocade SilkWorm 3200)</p> <p>EMC DS-8B3 (Brocade SilkWorm 3250)</p> <p>EMC DS-16B2 (Brocade SilkWorm 3800)</p> <p>EMC DS-16B3 (Brocade SilkWorm 3850)</p> <p>EMC DS-4100B (Brocade SilkWorm 4100)</p> <p>EMC DS-16M4 (McData Sphereon 4400)</p> <p>EMC DS-24M2 (McData Sphereon 4500)</p> <p>EMC DS-32M4 (McData Sphereon 4700)</p>
Fibre Channel switch firmware	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.

**Table 1-9. SAN-Attached Clusters Rules and Guidelines (continued)**

<b>Rule/Guideline</b>	<b>Description</b>
Fibre Channel HBAs supported	To determine which HBAs are supported on your PowerEdge system, see Table 1-2.
Fibre Channel HBA driver versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Fibre Channel HBA firmware versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
Dell EMC CX300, CX500, or CX700 core software	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.
EMC PowerPath versions	For a list of supported versions, see the "Dell Cluster Configuration Support Matrix" on page 24.

\* Applicable only to PowerEdge 1855, 1955, M600, and M605.

## Obtaining More Information

For more information about SAN-attached clusters, see the *Dell PowerEdge Cluster FE500W Systems Installation and Troubleshooting Guide*. You can also see the Dell Support website at [support.dell.com](http://support.dell.com).

## Dell Cluster Configuration Support Matrix

*The Dell Cluster Configuration Support Matrix* provides the latest supported drivers, firmware, and the operating system versions for your PowerEdge Cluster FE500W solution.

To obtain the supported drivers and firmware version for the operating system that is running on your PowerEdge Cluster FE500W cluster nodes:

- 1 Open a Web browser.
- 2 Navigate to the Dell High Availability Clustering website at [www.dell.com/ha](http://www.dell.com/ha).
- 3 Click the **Products & Services** tab.



- 4 Under **Product Offerings**, click **FE500W**.

The *Dell Cluster Configuration Support Matrix* for the PowerEdge Cluster FE500W appears.

- 5 Locate the appropriate drivers and firmware that are supported on your hardware and software components and the operating system version.

## Best Practices

This section provides best practices for troubleshooting issues that may occur in your cluster.

### Using a Tape Backup Library in a SAN

PowerEdge Cluster FE500W solutions running Windows 2000 Advanced Server that are configured with Emulex HBAs can be connected to one or more tape backup libraries that can be shared with the other clusters and systems in a SAN. To avoid disrupting I/O activities from other network systems to the tape drive and to ensure cluster failover operations, disable the target reset to the tape device.

To disable the target reset, perform the following steps:

- 1 Click the **Start** button, select **Run**, and type the following:

```
c:\Program Files\HBAnyware\elxcfg.exe --emc
```

The **Emulex Configuration Tool** window appears.

- 2 In the **Available Adapters** box, select the first HBA in the list.
- 3 In the **Adapter Controls** box, select **Disable Target Reset for Tape Devices**.
- 4 In the **File** menu, click **Apply**.
- 5 In the **Available Adapters** box, select the second HBA in the list.
- 6 Repeat step 3 and step 4.
- 7 Reboot the cluster node.
- 8 Repeat step 1 through step 7 on each additional node.

## **The Cluster Disks are not Initialized in Disk Management**

On clusters running Windows Server 2003, Disk Management may display the cluster disks as `not initialized`.

This issue may occur if the cluster disks are owned by other nodes in the same cluster. This behavior is normal and does not affect cluster operations.

For more information, see the Microsoft knowledge base article 818878 on the Microsoft Support website at [support.microsoft.com](http://support.microsoft.com).