



Dell® PowerEdge® Cluster (PowerEdge 4200)

INSTALLATION AND TROUBLESHOOTING GUIDE



Dell® PowerEdge® Cluster (PowerEdge 4200)

INSTALLATION AND TROUBLESHOOTING GUIDE



Information in this document is subject to change without notice.

© 1997 Dell Computer Corporation. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, and *PowerEdge* are registered trademarks and *DellWare* is a registered service mark of Dell Computer Corporation; *Intel*, *Pentium*, and *LANdesk* are registered trademarks of Intel Corporation; *Microsoft*, *Windows NT*, and *MS-DOS* are registered trademarks of Microsoft Corporation; *3Com* is a registered trademark of 3Com 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 Computer Corporation disclaims any proprietary interest in trademarks and trade names other than its own.

December 1997 P/N 17088

Safety Instructions

Use the following safety guidelines to help protect your computer system from potential damage and to ensure your own personal safety. See the Preface in this guide for information about the notational conventions used in this manual, including the distinction between warnings, cautions, and notes.

Before You Begin

Observe the following warnings while servicing this system:

WARNING: The power supplies in this computer system produce high voltages and energy hazards, which can cause bodily harm. Only trained service technicians are authorized to remove the computer covers and access any of the components inside the computer.

WARNING: This system may have more than one power supply cable. To reduce the risk of electrical shock, a trained service technician must disconnect all power supply cables before servicing the system.



ADVARSEL: Dette system kan have mere end et strømforsyningskabel. For at reducere risikoen for elektrisk stød, bør en professionel servicetekniker frakoble alle strømforsyningskabler, før systemet serviceres.

VAROITUS: Tässä järjestelmässä voi olla useampi kuin yksi virtajohto. Sähköiskuvaaran pienentämiseksi ammattitaitoisen huoltohenkilön on irrotettava kaikki virtajohdot ennen järjestelmän huoltamista.



OSTRZEŻENIE: System ten może mieć więcej niż jeden kabel zasilania. Aby zmniejszyć ryzyko porażenia prądem, przed naprawą lub konserwacją systemu wszystkie kable zasilania powinny być odłączone przez przeszkolonego technika obsługi.

ADVARSEL! Det er mulig at dette systemet har mer enn én strømledning. Unngå fare for stød: En erfaren servicetekniker må koble fra alle strømledninger før det utføres service på systemet.

WARNING: Detta system kan ha flera nätkablar. En behörig servicetekniker måste koppla loss alla nätkablar innan service utförs för att minska risken för elektriska stötar.

When Working Inside the Computer

Before taking the covers off of the computer, perform the following steps in the sequence indicated:

1. **Turn off the computer and any peripherals.**
2. **Disconnect the computer and peripherals from their power sources. Also, disconnect any telephone or telecommunications lines from the computer.**

Doing so reduces the potential for personal injury or shock.

3. **Touch an unpainted metal surface on the computer chassis, such as the power supply, before touching anything inside the computer.**

While you work, periodically touch an unpainted metal surface on the computer chassis to dissipate any static electricity that might harm internal components.

In addition, take note of these safety guidelines when appropriate:

- To help avoid possible damage to the system board, wait 5 seconds after turning off the system before removing a component from the system board or disconnecting a peripheral device from the computer.
- When you disconnect a cable, pull on its connector or on its strain-relief loop, not on the cable itself. Some cables have a connector with locking tabs; if you are disconnecting this type of cable, press in on the locking tabs before disconnecting the cable. As you pull connectors apart, keep them evenly aligned to avoid bending any connector pins. Also, before you connect a cable, make sure both connectors are correctly oriented and aligned.
- Handle components and cards with care. Don't touch the components or contacts on a card. Hold a card by its edges or by its metal mounting bracket. Hold a component such as a microprocessor chip by its edges, not by its pins.

WARNING

There is a danger of a new battery exploding if it is incorrectly installed. Replace the battery only with the same or equivalent type recommended by the manufacturer. Discard used batteries according to the manufacturer's instructions.

Protecting Against Electrostatic Discharge

Static electricity can harm delicate components inside the computer. To prevent static damage, discharge static electricity from your body before you touch any of the computer's electronic components, such as the microprocessor. You can do so by touching an unpainted metal surface on the computer chassis.

As you continue to work inside the computer, periodically touch an unpainted metal surface to remove any static charge your body may have accumulated.

You can also take the following steps to prevent damage from electrostatic discharge (ESD):

- When unpacking a static-sensitive component from its shipping carton, do not remove the component's antistatic packing material until you are ready to install the component in the computer. Just before unwrapping the antistatic packaging, be sure to discharge static electricity from your body.
- When transporting a sensitive component, first place it in an antistatic container or packaging.
- Handle all sensitive components in a static-safe area. If possible, use antistatic floor pads and workbench pads.

The following caution may appear throughout this document to remind you of these precautions:

CAUTION: See "Protecting Against Electrostatic Discharge" in the safety instructions at the front of this guide.

When Using the Computer System

As you use the computer system, observe the following safety guidelines:

- If your computer has a voltage selection switch on the power supply, be sure the switch is set to match the alternating current (AC) power available at your location:
 - 115 volts (V)/60 hertz (Hz) in most of North and South America and some Far Eastern countries such as Japan, South Korea, and Taiwan
 - 230 V/50 Hz in most of Europe, the Middle East, and the Far East
- Be sure the monitor and attached peripherals are electrically rated to operate with the AC power available in your location.
- To help prevent electric shock, plug the computer and peripheral power cables into properly grounded power sources. These cables are equipped with three-prong plugs to ensure proper grounding. Do not use adapter plugs or remove the grounding prong from a cable. If you must use an extension cable, use a three-wire cable with properly grounded plugs.
- To help protect the computer system from sudden, transient increases and decreases in electrical power, use a surge suppressor, line conditioner, or uninterruptible power supply (UPS).
- Be sure nothing rests on the computer system's cables and that the cables are not located where they can be stepped on or tripped over.
- Do not push any objects into the openings of the computer. Doing so can cause fire or electric shock by shorting out interior components.
- Keep the computer away from radiators and heat sources. Also, do not block cooling vents. Avoid placing loose papers underneath the computer, and do not place the computer in a closed-in wall unit or on a rug.

Preface

About This Guide

This guide provides information about installing, configuring, and troubleshooting the hardware and software components of the Dell PowerEdge Cluster. This document addresses the use of two PowerEdge 4200 server systems and one or two PowerEdge Scalable Disk System 100 (SDS 100) storage systems in the PowerEdge Cluster. Dell plans future clustering products that will incorporate other products in the Dell server family. User documentation specific to those systems will be available as new cluster products are released.

This guide addresses two audience levels:

- Users and system installers who will perform general setup, cabling, and configuration of the PowerEdge Cluster
- Trained service technicians who will perform more extensive installations such as firmware upgrades and installation of required expansion cards

This guide identifies the appropriate audience for each topic being discussed.

The chapters and appendixes in this guide are summarized as follows:

- Chapter 1, “Getting Started,” provides an overview of the PowerEdge Cluster and outlines the steps for installing a new PowerEdge Cluster system or modifying an existing PowerEdge system into a PowerEdge Cluster.
- Chapter 2, “Cabling the Cluster Hardware,” provides instructions for properly cabling the system hardware components.
- Chapter 3, “Configuring the Cluster Software,” describes the software configuration options that must be specified to properly set up the cluster system.
- Chapter 4, “Running Applications on a Cluster” provides general information about running applications on the PowerEdge Cluster.
- Chapter 5, “Troubleshooting,” provides information to help you troubleshoot problems with the cluster’s installation and configuration.
- Appendix A, “Upgrading to a Cluster Configuration,” provides specific information to service technicians about upgrading existing system hardware and software to a cluster configuration.
- Appendix B, “Stand-Alone and Rack Configurations,” lists the Dell-supported stand-alone and rack configurations and provides instructions for installing the network switch in a rack.
- Appendix C, “Cluster Data Sheet,” provides a form for gathering and recording important information about your PowerEdge Cluster.
- Appendix D, “PowerEdge Cluster Configuration Matrix,” describes the configuration matrix form, which is used to record information about the cluster hardware such as service tag numbers and types of adapters installed in the cluster node PCI slots.
- Appendix E, “Regulatory Compliance,” lists the regulatory standards with which the PowerEdge Cluster has been tested and certified for compliance.
- Appendix F, “Safety Information for Technicians,” provides important safety warnings about electrostatic discharge (ESD).

- Appendix G, “Warranties and Return Policy,” describes the warranty information pertaining to the system.

Other Documentation You May Need

You may need to reference the following documentation when performing the procedures in this guide:

- The *Dell PowerEdge 4200 Systems User’s Guide*, which describes system features and technical specifications, small computer system interface (SCSI) device drivers, the System Setup program, software support, and the system configuration utility.
- The *Dell PowerEdge SDS 100 Storage System Installation and Service Guide*, which provides installation and operation instructions for the PowerEdge SDS 100 storage system.
- The Intel LANDesk Server Manager software, which includes a CD containing the server manager software and the following documents: *LANDesk Server Manager Setup Guide*, *LANDesk Server Manager User’s Guide*, *LANDesk Server Control Installation and User’s Guide*, and *LANDesk Server Monitor Module Installation and User’s Guide*.
- The *Dell Hardware Instrumentation Package for Intel LANDesk Server Manager User’s Guide*, which provides installation and configuration procedures as well as the alert messages issued by this server management software.
- The *Using the Dell Server Assistant CD* document, which provides instructions for using the *Dell Server Assistant CD*.

You may also have one or more of the following documents:

- The *Dell PowerEdge Rack-Mountable Solutions Installation Guide*, *Dell PowerEdge 4xxx and 6xxx Systems Rack Kit Installation Guide*, and *Dell PowerEdge SDS 100 Storage System Rack Installation Guide*, which provide detailed instructions for installing the cluster components in a rack.
- The following documents accompany the Dell PowerEdge Expandable Redundant Array of Inexpensive Disks (RAID) Controller: *Dell PowerEdge*

Expandable RAID Controller User’s Guide, *Dell PowerEdge Expandable RAID Controller Client User’s Guide*, *Dell PowerEdge Expandable RAID Controller General Alert Server User’s Guide*, and *Dell PowerEdge Expandable RAID Controller Battery Backup Module User’s Guide*.

- Documentation for the Microsoft Windows NT Server Enterprise Edition operating system is included with the system if you ordered the operating system software from Dell. This documentation describes how to install (if necessary), configure, and use the operating system software.
- Documentation is included with any options you purchase separately from the system. This documentation includes information that you need to configure and install these options in the Dell computer.
- Technical information files—sometimes called “readme” files—may be installed on the hard-disk drive to provide last-minute updates about technical changes to the system or advanced technical reference material intended for experienced users or technicians.

*NOTE: Documentation updates are sometimes included with the system to describe changes to the system or software. Always read these updates **before** consulting any other documentation because the updates often contain information that supersedes the information in the other documents.*

Notational Conventions

The following subsections list notational conventions used in this document.

Warnings, Cautions, and Notes

Throughout this guide, there may be blocks of text printed in bold type within boxes or in italic type. These blocks are warnings, cautions, and notes, and they are used as follows:

WARNING: A WARNING indicates the potential for bodily harm and tells you how to avoid the problem.

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

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

Typographical Conventions

The following list defines (where appropriate) and illustrates typographical conventions used as visual cues for specific elements of text throughout this document:

- *Keycaps*, the labeling that appears on the keys on a keyboard, are enclosed in angle brackets.
Example: <Enter>
- *Key combinations* are a series of keys to be pressed simultaneously (unless otherwise indicated) to perform a single function.
Example: <Ctrl><Alt><Enter>
- *Commands* presented in lowercase bold are for reference purposes only and are not intended to be typed when referenced.
Example: “Use the **format** command to . . .”

In contrast, commands presented in the Courier New font are part of an instruction and intended to be typed.
Example: “Type `format a:` to format the diskette in drive A.”

- *Filenames* and *directory names* are presented in lowercase bold.
Examples: **autoexec.bat** and **c:\windows**
- *Syntax lines* consist of a command and all its possible parameters. Commands are displayed in lowercase bold; variable parameters (those for which you substitute a value) are displayed in lowercase italics; constant parameters are displayed in lowercase bold. The brackets indicate items that are optional.
Example: **del** [*drive:*] [*path*]*filename* [*/p*]
- *Command lines* consist of a command and may include one or more of the command’s possible parameters. Command lines are presented in the Courier New font.
Example: `del c:\myfile.doc`
- *Screen text* is text that appears on the screen of your monitor or display. It can be a system message, for example, or it can be text that you are instructed to type as part of a command (referred to as a *command line*). Screen text is presented in the Courier New font.
Example: The following message appears on your screen:
`No boot device available`
Example: “Type `md c:\dos` and press <Enter>.”
- *Variables* are placeholders for which you substitute a value. They are presented in italics.
Example: `SIMMn` (where *n* represents the SIMM socket designation)

Contents

Chapter 1

Getting Started 1-1

PowerEdge Cluster Components	1-1
Minimum System Requirements	1-2
Basic Installation Procedure	1-3
Adding Peripherals Required for Clustering	1-4
Setting Up the Cluster Hardware	1-5
Cabling the Cluster Hardware	1-5
Updating System BIOS/Firmware for Clustering	1-5
Setting Up the Shared Storage Subsystem Hard-Disk Drives	1-6
Setting Up the Internal SCSI Hard-Disk Drives	1-6
Installing and Configuring Windows NT Server Enterprise Edition	1-6
Installing and Configuring the Microsoft Cluster Server Software	1-6
Installing PowerEdge Cluster Applications	1-6
Checking the System	1-6

Chapter 2

Cabling the Cluster Hardware 2-1

Cluster Cabling	2-1
One Shared Storage Subsystem Cabled to a Cluster	2-1
Two SDS 100 Storage Systems Cabled to a Single RAID Controller	2-3
Two SDS 100 Storage Systems Cabled to Dual RAID Controllers	2-4
SMB Cabling	2-5
NIC Cabling	2-5
Power Cabling	2-7
Mouse, Keyboard, and Monitor Cabling	2-7
Disconnecting SCSI Cables While the Cluster Is Running	2-7

Chapter 3

Configuring the Cluster Software 3-1

Low-Level Software Configuration	3-1
Important System Warning	3-1
SCSI Host Adapter IDs	3-2
Disabling a RAID Controller BIOS	3-2
RAID Level for the Shared Storage Subsystem(s)	3-2
RAID Level for the Internal Hard-Disk Drives (Optional).	3-2
High-Level Software Configuration	3-3
Installing Intel LANDesk Server Manager	3-3
Choosing a Domain Model	3-3
Static IP Addresses	3-3
IPs and Subnet Masks	3-3
Configuring Separate Networks on a Cluster	3-3
Changing the IP Address of a Cluster Node	3-4
Naming and Formatting Shared Drives.	3-4
Driver for the RAID Controller.	3-4
Updating the NIC Driver.	3-5
Adjusting the Paging File Size and Registry Sizes	3-5
Verifying the Cluster Functionality	3-5
1 x 8 Mode on the SDS 100 Storage System	3-5
SCSI Controller IDs	3-6
Cluster Domain	3-6
RAID Controller Driver	3-6
Shared Storage Subsystem Drive Letters	3-6
Cluster Network Communications	3-6
Cluster Service	3-7
Availability of Cluster Resources	3-7
Uninstalling Microsoft Cluster Server	3-7
Removing a Node From a Cluster	3-7
Setting Up the Quorum Resource.	3-8
Using the ftdisk Driver	3-8
Cluster RAID Controller Functionality	3-8
Rebuild Function Does Not Complete After Reboot or Power Loss	3-8
Rebuild Rate Not Adjustable on Cluster-Enabled RAID Controller.	3-8
Using the Maximize Feature in PowerEdge RAID Console.	3-8
Rebuild Operation in RAID Console	3-9

Chapter 4	
Running Applications on a Cluster	4-1
Setting Up Applications	
Software to Run on the Cluster	4-1
Internet Information Server Service	4-1
File Share Service	4-2
Print Spooler Service	4-3
Using the Rediscovery Application in Intel LANDesk	4-4
Running chkdsk /f on a Quorum Disk	4-5
Tape Backup for Clustered Systems	4-5
Chapter 5	
Troubleshooting	5-1
Appendix A	
Upgrading to a Cluster Configuration	A-1
Checking Your Existing Hardware	A-1
Adding Expansion Cards for a Cluster Upgrade	A-1
Mounting, Cabling, and Configuring the Cluster Hardware	A-2
Installing and Configuring the Cluster Software	A-3
Upgrading the PowerEdge 4200 Firmware	A-3
Upgrading the PowerEdge SDS 100 Storage System Firmware	A-3
Setting the Cluster Mode With BIOS Setup	A-3
Installing and Configuring NICs	A-3
Appendix B	
Stand-Alone and Rack Configurations	B-1
Power Requirements of the PowerEdge Cluster	B-1
Supported Stand-Alone Configurations	B-2
Rack Safety Notices	B-2
Kit Installation Restrictions	B-2
Rack Stabilizer Feet	B-2
Supported Rack Configuration	B-5
Rack-Mounting the Network Switch	B-6
Appendix C	
Cluster Data Sheet	C-1
Appendix D	
PowerEdge Cluster Configuration Matrix	D-1

Appendix E

Regulatory Compliance. E-1

Regulatory Standards	E-1
CE Notice	E-1
Safety Standard	E-1
EMC Standards	E-1

Appendix F

Safety Information for Technicians F-1

Appendix G

Warranties and Return Policy. G-1

Limited Three-Year Warranty (U.S. and Canada Only)	G-1
Coverage During Year One	G-1
Coverage During Years Two and Three	G-2
General	G-2
“Total Satisfaction” Return Policy (U.S. and Canada Only)	G-2

Index

Figures

Figure 1-1. PowerEdge Cluster Layout	1-2
Figure 1-2. Back View of a PowerEdge 4200 Cluster Node	1-5
Figure 2-1. Cabling a Clustered System With One PowerEdge SDS 100 Storage System.	2-2
Figure 2-2. Cabling Single RAID Controllers to Two PowerEdge SDS 100 Storage Systems	2-3
Figure 2-3. Cabling Dual RAID Controllers to Two PowerEdge SDS 100 Storage Systems	2-4
Figure 2-4. SMB Cable Connected to One SDS 100 Storage System	2-5
Figure 2-5. SMB Cables Connected to Two SDS 100 Storage Systems	2-5
Figure 2-6. Cabling the Network Switch	2-6
Figure 2-7. PowerEdge Cluster Power Cabling	2-8
Figure A-1. Back View of a PowerEdge 4200 Cluster Node	A-2
Figure B-1. Supported Stand-Alone Configurations With One SDS 100 Storage System.	B-3
Figure B-2. Supported Stand-Alone Configurations With Two SDS 100 Storage Systems	B-4
Figure B-3. Supported Rack Configuration	B-5

Figure B-4. Attaching the Rack-Mounting Hardware on the Network Switch . . . B-6
Figure D-1. PowerEdge Cluster Configuration Matrix D-2

Table

Table 5-1. Troubleshooting 5-1

Chapter 1

Getting Started

The Dell® PowerEdge® Cluster is an enterprise system that implements clustering technology based on the Microsoft® Windows NT® Server, Enterprise Edition 4.0 operating system and Microsoft Windows NT Cluster Server. The Dell PowerEdge Cluster provides the following benefits in meeting the needs of mission-critical network applications:

- High availability of system services and resources to network clients
- Redundant storage of application data
- Failure recovery for cluster applications
- Capability to repair, maintain, or upgrade a cluster server without taking the whole cluster off-line
- Sharing of processing and communication work load between the two servers

The term *cluster* refers to two or more server systems (referred to as *nodes*) that are interconnected with appropriate hardware and software to provide a single point of continuous access to network services (for example, file service, database applications, resources, and so on) for network clients. Each cluster node is configured with software and network resources that enable it to interact with the other node to provide a mutual redundancy of operation and application processing. Because the servers interact in this way, they appear as a single system to the network clients.

As an integrated system, the PowerEdge Cluster is designed to handle most hardware failures and downtime

dynamically. In the event that one of the cluster nodes fails or experiences downtime, the processing work load of the failed node switches over (or *fails over*) to the remaining node in the cluster. This fail-over capability enables the cluster system to keep network resources and applications up and running on the network while the failed node is taken off-line, repaired, and brought back online. The overall impact of a node failure to network operation is minimal.

PowerEdge Cluster Components

The Dell PowerEdge Cluster consists of two Dell PowerEdge 4200 systems (the cluster nodes) equipped with one or two Dell PowerEdge Expandable redundant array of inexpensive disks (RAID) Controllers and two network interface controllers (NICs) to provide a dedicated node-to-node network interconnection and a regular Ethernet local area network (LAN) connection. Each server has shared Ultra/Wide small computer system interface (SCSI) connections to one or more Dell PowerEdge Scalable Disk System (SDS 100) storage system(s). Figure 1-1 shows a layout of the PowerEdge Cluster components and their interconnections.

Each component of the PowerEdge Cluster has a minimum system requirement. The following section lists and describes the minimum system requirements for the PowerEdge Cluster.

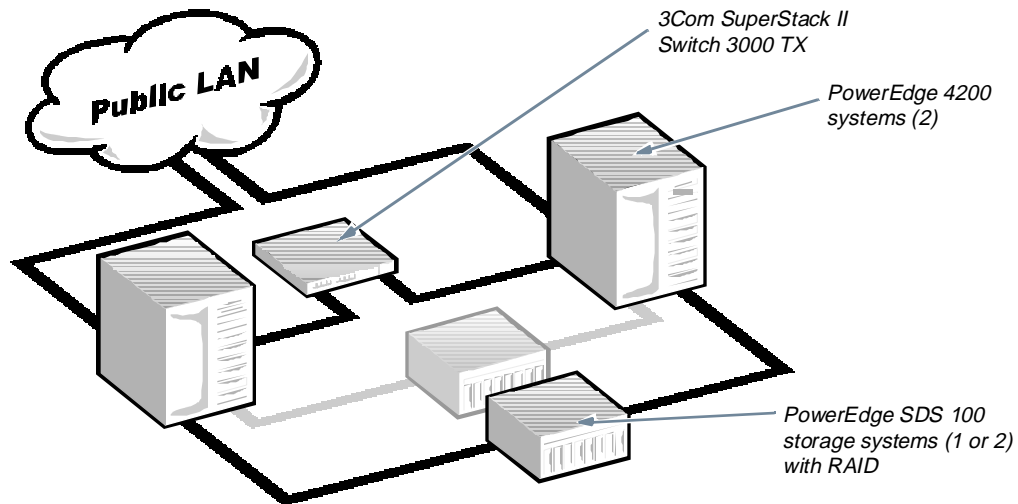


Figure 1-1. PowerEdge Cluster Layout

Minimum System Requirements

NOTE: If you are upgrading an existing system to a PowerEdge Cluster, check this list to ensure that your upgrade meets these requirements.

The PowerEdge Cluster requires the following minimum system hardware configuration:

- Two PowerEdge 4200 systems with the following configuration:
 - One or two 233-megahertz (MHz), one or two 266-MHz, or one or two 300-MHz Intel® Pentium® II microprocessors with at least 512 kilobytes (KB) of level 2 (L2) cache.
 - 128 megabytes (MB) of random-access memory (RAM).
 - A minimum of one PowerEdge Expandable RAID Controller in each PowerEdge system with 16 MB of single in-line memory module (SIMM) memory. This controller must have cluster-specific firmware and must be installed in Peripheral Component Interconnect [PCI] slot 7. A second cluster RAID controller can be added to slot 5, but the first cluster RAID controller must be installed in slot 7.
- Two 4-GB internal SCSI hard-disk drives (three drives are required for an internal RAID 5 configuration).
- Two Ethernet NICs, installed in PCI slots 4 and 8. The LAN-connected NIC resides in PCI slot 8, and the node-to-node interconnect NIC occupies slot 4.
- Power cabling and distribution components required:
 - For the Americas: Two Power Techniques power strips with Type B plugs, Model P905200
 - For Europe: One or two Marway power distribution units (PDUs), Model MPD-411013 or two Power Techniques power strips with Type B plugs, Model P906200
- One or two SDS 100 storage system(s) for the shared disk resource with the following configuration:
 - Cluster-specific basic input/output system (BIOS) upgrade for the PowerEdge systems for turning the SDS 100 storage system backplane into a 1 x 8 mode (one SCSI channel with up to eight hard-disk drives) when two RAID controllers are present.

- At least three SCSI hard-disk drives in each SDS 100 storage system to support RAID 5 functionality. Microsoft Cluster Server currently supports only the Microsoft Windows NT file system (NTFS) format for the shared storage subsystem.
- Two 4-meter (m) SCSI cables for each SDS 100 storage system in the cluster.
- A 3Com[®] SuperStack II Switch 3000 TX 8-port switch and accessories, which includes the following:
 - Four Category 5 unshielded twisted pair (UTP) Ethernet cables
 - Hardware for mounting the network switch in a Dell Rack-Mountable Solutions enclosure (optional)

In addition to the preceding hardware components, the following software components are also required:

- Windows NT Server, Enterprise Edition 4.0 operating system installed on the PowerEdge systems. Two Windows NT Server Enterprise Edition licenses are required, plus workstation licenses for all the client systems running on the network.
- Clustering software recovery kits for the customer environment. These recovery kits are in addition to the standard file, print, and Internet Information Server (IIS) resources that are bundled with the Cluster Server software.
- Transmission Control Protocol/Internet Protocol (TCP/IP) running on the LAN. NetBIOS Extended User Interface (NetBEUI) and Internet Packet Exchange/Sequenced Packet Exchange (IPX/SPX) are not supported.
- Server Management Agent rediscovery application.
- Cluster-specific Windows NT Server driver for the PowerEdge Expandable RAID Controllers.

Cluster requires two alternating current (AC) circuits with a minimum load capacity of 20 amperes each.

Your installation of the PowerEdge Cluster may be either a completely new installation or an upgrade of an existing system. If your PowerEdge Cluster is completely new, the operating system and some applications may be installed on your system already. Installation in this case is a matter of setting up and cabling the hardware, setting some configuration options, setting network addresses, and performing checks on the system.

If you are upgrading existing equipment, several additional steps must be performed such as installing additional NIC and RAID expansion cards, updating firmware, and installing both the operating system and cluster software on each cluster node. Hardware installation and updates to firmware should be performed only by trained service technicians.

The following is a comprehensive list of the steps that may be required to install a PowerEdge Cluster, whether it is a new system installation or an upgrade to an existing system:

- 1. For system upgrades, add NICs, RAID controllers, hard-disk drives, and so on to the existing system hardware to meet the requirements for a clustered system.**
- 2. Set up the cluster equipment in either a stand-alone or rack configuration.**
- 3. Cable the system hardware for clustering.**
- 4. For system upgrades, update the existing system components with cluster-specific firmware.**
- 5. If not already done, configure the RAID level on the shared storage subsystem using the PowerEdge Expandable RAID Controller BIOS configuration utility.**
- 6. If not already done, partition and format the hard-disk drives in the shared storage subsystem(s). Also, partition and format any new hard-disk drives added to the cluster nodes for a system upgrade.**
- 7. If not already done, install and/or configure Windows NT Server Enterprise Edition and the included Service Pack on each cluster node.**

Basic Installation Procedure

NOTE: Before installing the PowerEdge Cluster, ensure that your site power is adequate to handle the power requirements of the cluster equipment. PowerEdge

8. **Configure the public and private NIC interconnects in each node and place the interconnects on separate IP subnetworks.**
9. **If not already done, install and/or configure the Microsoft Cluster Server software on each cluster node.**
10. **Check out the functionality of the fully installed cluster.**
11. **Install and set up applications.**

The following sections briefly describe each of these steps.

Adding Peripherals Required for Clustering

NOTE: Hardware installation should be performed only by trained service technicians.

If you are upgrading your existing hardware to a cluster configuration, additional peripheral devices and expansion cards need to be added to the system to meet the minimum cluster requirements listed earlier in this chapter.

For example, you need to install a second NIC card to ensure that the system meets the minimum configuration of two NIC cards—one card in PCI slot 8 for the public LAN connection and another card in PCI slot 4 for the node-to-node interconnection. You also need to add a cluster-enabled PowerEdge Expandable RAID Controller for the required shared storage subsystem used by the two nodes.

NOTE: The first cluster-enabled PowerEdge Expandable RAID Controller must be installed in PCI slot 7.

Figure 1-2 shows the placement of these devices in a PowerEdge 4200 system. See Appendix A for further information about upgrading an existing PowerEdge 4200 system with expansion cards required for clustering.

Additionally, you may need to add hard-disk drives and another PowerEdge Expandable RAID Controller to the PowerEdge system if you are configuring the system's internal drives as a RAID array. However, this is not a requirement for clustering. Refer to the *Dell PowerEdge 4200 Systems Installation and Troubleshooting Guide* for instructions on installing expansion cards or hard-disk drives in the PowerEdge 4200 system.

If you are upgrading an existing SDS 100 storage system to meet the cluster requirements for the shared storage subsystem, you may need to install additional hard-disk drives to the shared storage subsystem. The size and number of drives you add depend on the RAID level you want to use and the number of hard-disk drives already present in your system. For information on installing hard-disk drives in the PowerEdge SDS 100 storage system, refer to the *Dell PowerEdge SDS 100 Storage System Installation and Service Guide*.




Figure 1-2. Back View of a PowerEdge 4200 Cluster Node

Setting Up the Cluster Hardware

The PowerEdge Cluster can be set up in either a free-standing configuration or installed in a Dell Rack-Mountable Solutions enclosure. Information on Dell-supported rack configurations for the cluster is provided in Appendix B, “Stand-Alone and Rack Configurations,” in this guide. Also included in Appendix B are instructions for installing the network switch in a rack. For instructions on installing all other PowerEdge Cluster components (including the Apex Outlook Concentrator switch) in a Dell rack, refer to the *Dell PowerEdge Rack-Mountable Solutions Installation Guide*.

Cabling the Cluster Hardware

After the PowerEdge Cluster hardware is set up, the system must be properly cabled for clustering. Chapter 2, “Cabling the Cluster Hardware,” provides instructions for cabling the cluster components.

Updating System BIOS/Firmware for Clustering

NOTE: BIOS upgrades should be performed only when instructed by a Dell support technician.

If you are upgrading existing hardware to a PowerEdge Cluster, the BIOS firmware for the PowerEdge system(s) and SDS 100 storage system(s) must be updated to support clustering. Appendix A, “Upgrading to a Cluster Configuration,” provides instructions on performing all necessary firmware updates.

Setting Up the Shared Storage Subsystem Hard-Disk Drives

If your Power-Edge Cluster consists of all new components, the hard-disk drives in the shared storage subsystem may already be partitioned, formatted, and set up in a RAID configuration for clustering. If you are upgrading a shared storage subsystem in an existing system, the shared hard-disk drives must be set up for clustering as part of the upgrade.

The first step is to configure the RAID level that you will be using in your cluster. For instructions on setting up a RAID array, refer to the *Dell PowerEdge Expandable RAID Controller User's Guide*. Then the hard-disk drives in the shared storage subsystem must be partitioned and formatted, and drive letters must be assigned to each drive. For instructions on partitioning and formatting the shared storage subsystem hard-disk drives, refer to the *Microsoft Windows NT Server Enterprise Edition Administrator's Guide and Release Notes* and the *Dell PowerEdge Expandable RAID Controller User's Guide*. Chapter 3, "Configuring the Cluster Software," in this guide describes how to assign drive letters to the shared hard-disk drives.

Setting Up the Internal SCSI Hard-Disk Drives

If you have added new hard-disk drives to your Power-Edge system or are setting up the internal drives in a RAID configuration, you must configure the RAID, if applicable, and partition and format the drives before you can install Windows NT Server Enterprise Edition. For instructions on partitioning and formatting SCSI hard-disk drives, refer to your PowerEdge system *User's Guide*. For instructions on setting up a RAID array, refer to the *Dell PowerEdge Expandable RAID Controller User's Guide*.

Installing and Configuring Windows NT Server Enterprise Edition

If it has not already been done, Windows NT Server Enterprise Edition must be installed on the internal hard-disk drives of both cluster nodes.

NOTE: Windows NT Server Enterprise Edition cannot be run from the shared storage subsystem.

Cluster-specific device drivers are also installed at this time. Refer to the *Microsoft Windows NT Server Enterprise Edition Administrator's Guide and Release Notes* for instructions on installing and configuring the operating system and adding cluster-specific device drivers.

Refer to Chapter 3, "Configuring the Cluster Software," of this guide for information specific to configuring Windows NT Server Enterprise Edition on your PowerEdge Cluster.

Installing and Configuring the Microsoft Cluster Server Software

Like Windows NT Server Enterprise Edition, the Cluster Server software must be installed on both cluster nodes if it has not already been done. Refer to the *Microsoft Windows NT Cluster Server Administrator's Guide* for instructions on installing and configuring the clustering software. Also refer to Chapter 3, "Configuring the Cluster Software," for specific information about installing and configuring Microsoft Cluster Server on your PowerEdge Cluster.

Installing PowerEdge Cluster Applications

Additional steps are required to configure applications software to run on the cluster. Chapter 4 in this guide provides general information about this process and cites example procedures for setting up the Windows NT Internet Information Server (IIS) Virtual Root service, the File Share service, and the Print Spooler service to run on a cluster.

Chapter 4 also describes the rediscovery application, which must be run whenever the primary cluster node fails over to the secondary cluster node. The rediscovery application enables the secondary cluster node to rediscover and reestablish system management of the SDS 100 storage system(s).

Checking the System

When installation is complete, you should check the functionality of your cluster system by performing a number of tests. See "Verifying the Cluster Functionality" in Chapter 3 for specific tests and procedures that you can perform to check out the cluster.

Chapter 2

Cabling the Cluster Hardware

The Dell PowerEdge Cluster must be installed and cabled correctly to ensure that the cluster functions properly. This chapter instructs you on how to cable your system hardware for a cluster configuration.

Information about configuring your PowerEdge Cluster is provided in Chapter 3, “Configuring the Cluster Software.”

For instructions on installing the Microsoft Windows NT Server Enterprise Edition operating system and the Microsoft clustering software, refer to the *Microsoft Windows NT Server Enterprise Edition Administrator's Guide and Release Notes* and the *Microsoft Windows NT Cluster Server Administrator's Guide*. For installation and configuration information specific to the Dell PowerEdge 4200 systems or the Dell PowerEdge Scalable Disk System 100 (SDS 100) storage system, refer to the Dell documentation for those systems.

Cluster Cabling

The PowerEdge Cluster consists of two PowerEdge 4200 server systems, one or two PowerEdge SDS 100 storage systems, a 3Com SuperStack II Switch 3000 TX, and a pair of power strips or a single power distribution unit, depending on how the system is configured. These components are interconnected as follows:

- A 4-meter (m) small computer system interface (SCSI) cable is connected from the RAID controller in each PowerEdge system to the SDS 100 storage system(s).
- A system management bus (SMB) cable is connected from the SMB connector on one of the two PowerEdge systems (preferably the system designated as Node 1 or the primary node) to the SMB connector on the SDS 100 storage system.

- Category 5 Ethernet cables are connected from each of the network interface controllers (NICs) in each PowerEdge system to the 3Com switch.
- Power cables are connected according to the safety requirements for your region:
 - For customers in the Americas: Power cables for the cluster components are routed through two Power Techniques power strips. The primary power supplies of the cluster components are all cabled to one power strip, and the redundant power supplies on the components are all cabled to the second power strip. Each power strip is connected, via Type B plugs and connectors, to a separate alternating current (AC) circuit, each with a minimum power capacity of 20 amperes (amps).
 - For customers in Europe: All power cables are connected to one or two Marway power distribution units (PDUs), Model MPD-411013 or two Power Techniques power strips with Type B plugs, Model P906200.

The following sections describe each of these cabling procedures.

One Shared Storage Subsystem Cabled to a Cluster

Use the following procedure to connect your cluster system to a single SDS 100 storage system. Refer to Figure 2-1 for a diagram of the cabling scheme.

CAUTION: Do not turn on the PowerEdge 4200 systems or the SDS 100 storage system(s) until all cabling is complete.




Figure 2-1. Cabling a Clustered System With One PowerEdge SDS 100 Storage System

1. Connect the 68-pin connector on the 4-m SCSI cable to SCSI connector A on the back of the SDS 100 storage system, and tighten the retaining screws.
2. Connect the ultra-high density (UHD) connector of the SCSI cable to the channel 0 connector (the rightmost connector) on the cluster RAID controller in the first PowerEdge server, and tighten the retaining screws.
3. Connect the 68-pin connector on the second 4-m SCSI cable to SCSI connector B on the back of the SDS 100 storage system, and tighten the retaining screws.
4. Connect the UHD connector of the second SCSI cable to the channel 0 connector on the cluster RAID controller in the second PowerEdge server, and tighten the retaining screws.

NOTES: On clusters with a single SDS 100 storage system, either server system can be connected to either storage system SCSI connector.

Be sure to securely tighten the retaining screws on the SCSI connectors to ensure a reliable connection.

NOTE: If the SDS 100 storage system is ever disconnected from the cluster, it must be reconnected to the same controller channels on the RAID controllers to operate properly.

Two SDS 100 Storage Systems Cabled to a Single RAID Controller

Connecting the cluster to two SDS 100 storage systems is similar to connecting to a single SDS 100 storage system.

First, connect the channel 0 connector of each PowerEdge Expandable RAID Controller to the back of the first storage system, as described in the preceding section. Then connect the channel 1 connector of each RAID controller to the second storage system (see Figure 2-2).

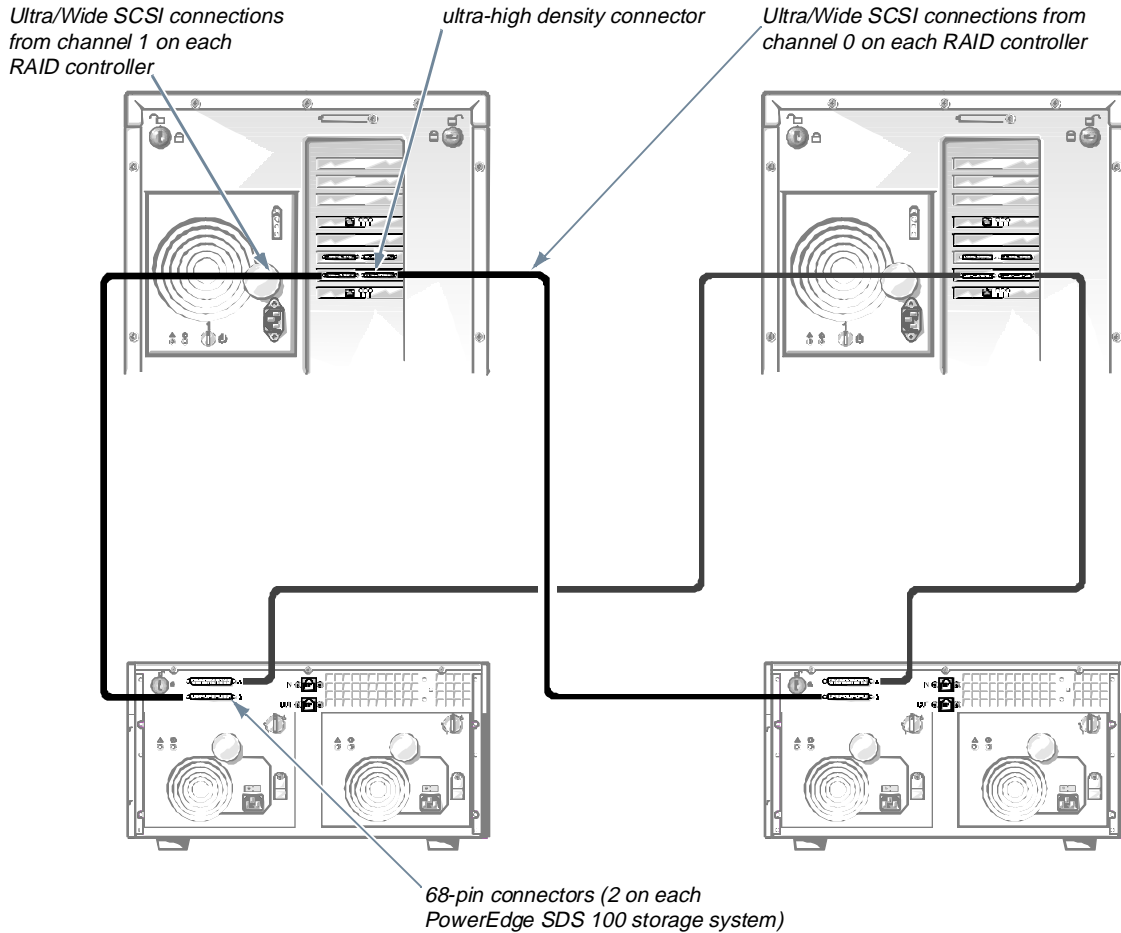


Figure 2-2. Cabling Single RAID Controllers to Two PowerEdge SDS 100 Storage Systems

Two SDS 100 Storage Systems Cabled to Dual RAID Controllers

To cable cluster nodes with dual RAID controllers to two SDS 100 storage systems, connect the channel 0 connector of each RAID controller of the primary node (or the first node) to the “A” connectors on the back of each storage system, and connect the channel 0 connectors of the secondary node’s RAID controllers to the “B” connectors on each storage system (see Figure 2-3).

NOTE: On clusters with multiple SDS 100 storage systems, the channel 0 connectors of the two RAID controllers must be connected to one storage system, and the channel 1 connectors must be connected to the second storage system. If the connections are ever removed, you must reconnect the cables as they were connected previously. To help ensure that the same storage system is attached to the same channels, tagging or color-coding the cables is recommended.




Figure 2-3. Cabling Dual RAID Controllers to Two PowerEdge SDS 100 Storage Systems

SMB Cabling

The SMB connector enables a host PowerEdge system to provide system-level management of the storage system(s).

NOTE: The SDS 100 storage system is connected to only one of the two PowerEdge systems in the cluster.

To install the SMB cable, use the following procedure:

1. **Connect one end of the SMB cable (supplied with the storage system) to the SMB connector labeled “IN” on the storage system’s back panel.**

Both connectors on the SMB cable are identical. The connectors are keyed for proper insertion.

2. **Connect the other end of the SMB cable to the SMB connector on the first PowerEdge system or to the SMB connector of the first storage system.**

- *If you are connecting only one storage system to the cluster, connect the SMB cable to the SMB connector on the first node of the cluster (see Figure 2-4).*




Figure 2-4. SMB Cable Connected to One SDS 100 Storage System

- *If you are connecting two SDS 100 storage systems to the cluster, link the storage systems in daisy-chain fashion to the PowerEdge system (see Figure 2-5). The first storage system in the*

chain connects to the SMB connector on the PowerEdge system’s back panel. The second storage system connects the SMB cable from its connector labeled “IN” to the connector labeled “OUT” on the first storage system’s back panel.




Figure 2-5. SMB Cables Connected to Two SDS 100 Storage Systems

NIC Cabling

The NICs in the PowerEdge systems provide two network connections on each node—a dedicated network interconnection between the cluster nodes and a connection to the local area network (LAN). Having two network interconnections from each PowerEdge system can provide redundancy at the communications level in case one of the cluster NICs fails.

The 3Com SuperStack II switch has eight ports available on its front panel, running at a switched rate of 100 megabits per second (Mbps). All ports on the SuperStack II switch are functionally identical, so the NIC cables can be attached to any of the ports in any order. Category 5 unshielded twisted-pair (UTP) cables are provided.

Figure 2-6 shows a sample configuration of NIC cabling where the private node-to-node interconnect (the NICs in PCI slot 4 of each node) routes through the network switch, and the LAN NICs connect directly to the public LAN. Other configurations are possible, including connecting all four NICs to the SuperStack II switch; however, in this scenario, the switch is a possible single point of failure.




Figure 2-6. Cabling the Network Switch

Power Cabling

Observe the following warnings when connecting the power cables to your PowerEdge Cluster system:

WARNING: Although each component of the PowerEdge Cluster meets leakage current safety requirements, the total leakage current may exceed the maximum that is permitted when the components are used together. To meet safety requirements in the Americas, you must use a Type B plug and socket connection for the cluster power to enable the appropriate level of ground protection. In Europe, you must use one or two power distribution units (PDUs) or two Type B plug and socket connections wired and installed by a qualified electrician in accordance with the local wiring regulations.

WARNING: Do not attempt to cable the PowerEdge Cluster to electrical power without first planning the distribution of the cluster's electrical load across available circuits. For operation in the Americas, the PowerEdge Cluster requires two AC circuits with a minimum capacity of 20 amps each to handle the electrical load of the system. Do not allow the electrical load of the system to exceed 16 amps on either circuit. For operation in Europe, the PowerEdge Cluster requires two circuits rated in excess of the combined load of the attached systems. Please refer to the ratings marked on the back of each cluster component when determining the total system's electrical load.

Figure 2-7 illustrates the proper power cabling of the PowerEdge Cluster components. Each component of the cluster must have power supplied by two separate AC circuits—one circuit to each component power supply. Therefore, the primary power supplies of all the PowerEdge Cluster components are grouped onto one circuit and the redundant power supplies are grouped onto another circuit.

Mouse, Keyboard, and Monitor Cabling

If you are installing the PowerEdge Cluster in a Dell Rack-Mountable Solutions cabinet, refer to the *Dell PowerEdge Rack-Mountable Solutions Installation Guide* for instructions on cabling each cluster node's mouse, keyboard, and monitor to the Apex Outlook switch box in the rack. The switch box enables you to use a single mouse, keyboard, and monitor for both systems.

Disconnecting SCSI Cables While the Cluster Is Running

If you must disconnect a SCSI cable between a powered-down server and a running SDS 100 storage system, you should first disconnect the cable from the back of the SDS 100 and then disconnect the cable from the RAID controller connector on the cluster node. This helps maintain the integrity of the SCSI signals while removing the cable.

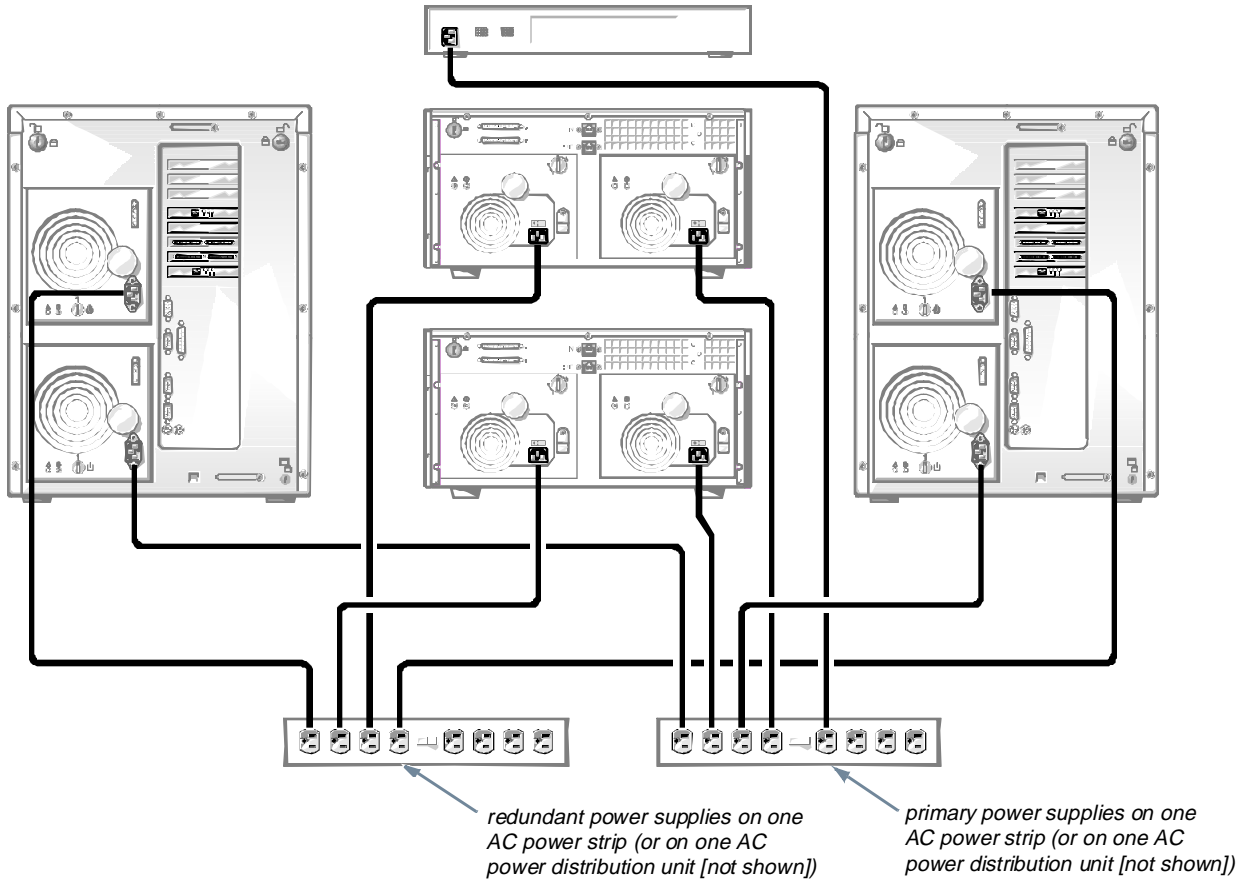


Figure 2-7. PowerEdge Cluster Power Cabling

Chapter 3

Configuring the Cluster Software

This chapter provides information about configuring the Dell PowerEdge Cluster system software. This guide does not provide instructions for installing the operating system or the cluster software. Installation instructions for the operating system are documented in the *Microsoft Windows NT Server Enterprise Edition Administrator's Guide and Release Notes*. Instructions for installing the Microsoft clustering software are provided in the *Microsoft Windows NT Cluster Server Administrator's Guide*. The information presented in this chapter serves as an addendum to the Microsoft documentation.

Before installing the Windows NT Server Enterprise Edition operating system or the Cluster Server software, you should have your system hardware properly cabled for clustering. See Chapter 2 in this guide for instructions on connecting the components of your PowerEdge Cluster.

Low-Level Software Configuration

Prior to installing Windows NT Server Enterprise Edition, you must make specific low-level software configurations to the PowerEdge Cluster. Low-level software configurations are settings you make to the system before the operating system is installed.

The following subsections describe the low-level software settings that must be made to your system to enable clustering.

Important System Warning

The following warning message appears on your screen whenever you attempt to modify the configuration of the shared storage subsystem on your cluster using either the PowerEdge Expandable RAID Controller BIOS configuration utility or the PowerEdge RAID Console utility:

```
!!!STOP!!!
```

This operation may change the configuration of disks and can cause loss of data!

Ensure:

1. Peer server is powered up for its controller NVRAM to be updated. Otherwise, disk configuration should be read from disk and saved to controller's NVRAM.
2. The second server must not be configuring the disks.
3. There is no I/O activity against shared drives.

The warning appears immediately when you activate the redundant array of inexpensive disks (RAID) basic input/output system (BIOS) configuration utility (by pressing <Ctrl><m> during the system's power-on self-test [POST]), or whenever you attempt to perform a data-destructive operation in the PowerEdge RAID Console utility. Examples of data-destructive operations include clearing the configuration of the logical drives or changing the RAID level of your shared hard-disk drives.

This warning alerts you to the possibility of data loss if certain precautions are not taken to protect the integrity of the data on your cluster. To prevent the loss of data, be sure that your cluster meets the following conditions before you attempt any data-destructive operation on your shared hard-disk drives:

- Be sure the peer server is powered up during the operation so that its RAID controller nonvolatile random-access memory (NVRAM) can be updated with the new configuration information. Alternately, if the peer server is down, you must save the disk configuration to the shared storage subsystem. When you restart the system later, update the peer server's NVRAM from the disk configuration saved to the shared storage subsystem.
- Be sure the peer cluster node is not currently configuring the shared storage subsystem.
- Be sure that no input/output (I/O) activity occurs on the shared storage subsystem during the operation.

SCSI Host Adapter IDs

On a small computer system interface (SCSI) bus, each device must have a unique SCSI identification (ID) number. The default SCSI ID of the RAID controller is 7. However, with RAID controllers from two cluster nodes occupying the same bus, the controller in the second cluster node must be set to a different SCSI ID to avoid a device conflict. Therefore, the RAID controller in the second cluster node should be set to SCSI ID 10. In addition, because multiple RAID controllers can reside on each node, all RAID controllers on the second node must be set specifically to SCSI ID 10.

Use the RAID controller BIOS configuration utility to set a SCSI ID. Start the utility by pressing <Ctrl><m> during the system's POST. From the Management Menu, select Objects, then select Adapter, then select the appropriate adapter (if applicable), and then select Initiator ID. If you are running the utility from the first cluster node, the SCSI ID should be set to 7. If you are on the second cluster node, change the 7 to 10 and press <Enter>. At the confirmation prompt, select Yes and then reboot the cluster node by pressing <Ctrl><Alt><Delete>.

Disabling a RAID Controller BIOS

The BIOS for all of the cluster-specific RAID controllers must be disabled. Only a RAID controller that is controlling the system boot device should have its BIOS enabled.

Use the RAID controller BIOS configuration utility to disable a cluster RAID. Start the utility by pressing <Ctrl><m> during POST. From the Management Menu, select Objects, then select Adapter, then select the appropriate adapter (if applicable), and then select Disable BIOS. Select the Disable BIOS setting if it is listed. If Enable BIOS is the choice that is offered, the BIOS for the RAID controller is already disabled.

RAID Level for the Shared Storage Subsystem(s)

The RAID level can be set using the RAID controller BIOS configuration utility. Start the utility by pressing <Ctrl><m> during POST. The recommended default RAID level for a cluster with two Dell PowerEdge Scalable Disk System 100 (SDS 100) storage systems is RAID 1+0. RAID 1+0 is a combination of RAID levels 1 and 0. Data is striped across the SDS 100 drives as in RAID 0. Each drive is mirrored on the second SDS 100, as in RAID 1. RAID 1+0 allows high-availability of the quorum resource, which can be mirrored on hard-disk drives on both SDS 100 systems.

For cluster systems with a single SDS 100 storage system, the recommended configuration consists of two logical drives—two of the SDS 100's hard-disk drives comprising the first logical drive and the remaining drives (up to six) comprising the second logical drive. The first logical drive should be configured for RAID 1 (disk mirroring) and should contain the quorum resource. The second logical drive should be configured for RAID 5 and should contain application data for the cluster.

RAID Level for the Internal Hard-Disk Drives (Optional)

Like the RAID level for the shared storage subsystem, this configuration can also be set using the RAID controller configuration utility. The recommended default configuration of the internal drives is a RAID 5. Additionally, the default channel for connecting the controller to the internal drives is channel zero.

High-Level Software Configuration

When the SCSI drives and RAID levels have been set up, Windows NT Server Enterprise Edition can be installed and configured. A number of operating system configurations must be set during the installation to enable clustering. These configuration requirements are described in the *Microsoft Windows NT Server Enterprise Edition Administrator's Guide and Release Notes*. The following subsections briefly discuss these configurations.

Installing Intel LANDesk® Server Manager

After installing the Windows NT Enterprise Edition operating system, install LANDesk prior to applying the Service Pack to your system. Refer to the *LANDesk Server Manager Setup Guide* for installation instructions.

Choosing a Domain Model

Cluster nodes can be set up in three possible configurations: as two stand-alone member servers, as two backup domain controllers (BDC), or as a primary domain controller (PDC) and a BDC. The first two configurations require an existing domain for the servers to join. The PDC/BDC configuration establishes a new domain of which the one server is the primary domain controller and the other server is the backup domain controller. Any of the three configurations can be chosen for clustering, but the recommended default is having each cluster server as a member server in an existing domain. This relieves the cluster nodes from the processing overhead involved in authenticating the user logon.

Static IP Addresses

The Microsoft Cluster Server software requires one static Internet Protocol (IP) address for the cluster and one static IP address for each disk resource group. A static IP address is an Internet address that a network administrator assigns exclusively to a system or a resource. The address assignment remains in effect until the network administrator changes it.

IPs and Subnet Masks

For the node-to-node network interface controller (NIC) connection on the PowerEdge Cluster, the default IP address 10.0.0.1 is assigned to the first node and the second node is assigned the default address 10.0.0.2. The default subnet mask is 255.0.0.0.

Configuring Separate Networks on a Cluster

Two network interconnects are strongly recommended for a cluster configuration to eliminate any single point of failure that could disrupt intracluster communication.

Separate networks can be configured on a cluster by redefining the network segment of the IP address assigned to the NICs residing in the cluster nodes.

For example, two NICs reside in two cluster nodes. The NICs in the first node have the following IP addresses and configuration:

NIC1:

IP address: 143.166.110.2
Default gateway: 143.166.111.3

NIC2

IP address: 143.166.111.3
Default gateway: 143.166.110.2

The NICs in the second node have the following IP addresses and configuration:

NIC1:

IP address: 143.166.110.4
Default gateway: 143.166.111.5

NIC2

IP address: 143.166.111.5
Default gateway: 143.166.110.4

IP routing is enabled and the subnet mask is 255.255.255.0 on all NICs.

The NIC1s of two machines establish one network segment, and the NIC2s create another. In each system, one NIC is defined to be the default gateway for the other NIC.

When a packet gets sent across the network from a local client, the source and destination IP addresses of the packet are inserted in the IP header. The system checks whether the network ID of the destination address matches the network ID of the source address. If they match, the packet is sent directly to the destination computer on the local network. If the network IDs do not match, the packet is forwarded to the default gateway for delivery.

Changing the IP Address of a Cluster Node

NOTE: To change the IP address of a cluster node, the Cluster Service running on that node must be stopped. Once the service is stopped, the IP address can be re-assigned and the server restarted.

While the node is down, the Cluster Administrator utility running on the second node indicates that the first node is down by showing its icon in red. When the node is restarted, the two nodes reestablish their connection and the Cluster Administrator changes the node icon back to blue to show that the node is back online.

Naming and Formatting Shared Drives

The logical drives of the shared storage subsystem must be assigned drive letters and then formatted as Windows NT file system (NTFS) drives. The assigned drive letters must be identical on both cluster nodes.

NOTE: Because the number of drive letters required by individual servers in a cluster may vary, it is recommended that the shared drives be named in reverse alphabetical order beginning with the letter “z.”

Use the following procedure to assign drive letters and format drives:

1. **Click the Start button, point to Programs, point to Administrative Tools (Common), and click Disk Administrator.**
2. **At the confirmation dialog box, click Yes to enter a signature on all new physical or logical drives.**
3. **Find the disk icon for the first unnamed, unformatted drive, right-click the icon, and select Create from the submenu.**
4. **In the dialog box, create a partition the size of the entire drive (the default setting) and click OK.**
5. **Click Yes to confirm the partition.**
6. **With the pointer on the same icon, right-click and select Assign Drive Letter from the submenu.**
7. **Type the letter you want to assign the drive (for example, z) and click OK.**
8. **Highlight and right-click the drive icon again and select Commit Changes Now from the submenu.**
9. **Click Yes to save the changes.**
10. **Click Yes to confirm that changes were made.**
11. **Right-click the drive icon again and select Format from the submenu.**
12. **At the dialog box, change the file system to NTFS, click Quick Format, and click Start.**

The NTFS file system format is required for shared-disk resources under Microsoft Cluster Server.
13. **Click OK at the warning.**
14. **Click OK to acknowledge that format is complete.**
15. **Click Close to close the dialog box.**
16. **Repeat steps 3 through 15 for each remaining drive.**
17. **Close the Disk Administrator dialog box.**

When all drives have been assigned drive letters and formatted, the identical drive letters for the shared drives must be assigned on the second cluster node. To do this, enter the Disk Administrator on the second cluster node, right-click each drive, and assign the same drive letter to each drive that was assigned on the first cluster node.

Driver for the RAID Controller

The RAID controller driver, **pedge.sys**, must be version 2.04 or later. Refer to the section in this chapter entitled “RAID Controller Driver” for instructions on how to verify that this driver is installed.

Updating the NIC Driver

Dell recommends that you use Windows NT NIC driver version 2.22 or later for the Intel Pro100B network controller. Perform the following procedure on both cluster nodes to update the NIC driver:

1. **Go to the Control Panel, double-click the Network icon, and click on the Adapters tab.**
2. **Highlight one of the adapters and click Update.**
3. **In the dialog box, type A : . Place the diskette containing the updated Intel Pro100B driver into drive A and press <Enter>.**

Windows NT installs the NIC driver.

4. **When the driver has been installed, click Close to exit the Network dialog box.**

Adjusting the Paging File Size and Registry Sizes

To enable adequate system resources for clustering, it is recommended that you increase the paging file and registry file sizes on the cluster nodes. Set the paging file size to at least twice the capacity of the system RAM, up to 256 megabytes (MB). For systems with RAM capacities over 256 MB, set the paging file size at or above the capacity of the RAM up to the available free space on your hard-disk drive. Set the registry file size to at least 64 MB. These adjustments can be made prior to applying the current Windows NT Service Pack.

Use the following procedure to make the paging file and registry size adjustments on each cluster node:

1. **Go to Control Panel, double-click the System icon, and click the Performance tab to see the System Properties dialog box. In the Virtual Memory group box, click Change.**
2. **In the dialog box, set the Paging File maximum size to 256 MB. Set the Registry File size to 64 MB and click OK.**
3. **When asked to restart the system, click No.**
4. **Apply the current Service Pack for Windows NT Server Enterprise Edition.**

Verifying the Cluster Functionality

To ensure that the PowerEdge Cluster functions properly, you should perform a series of checks of the system's operation and configuration. These checks should be performed to verify that the cluster meets the following conditions:

- Each SDS 100 storage system is running in 1 x 8 mode.
- The controller IDs on each shared bus are different.
- All cluster servers and clients are able to log on to the same domain.
- The cluster-specific driver for the RAID controller is installed on both cluster nodes.
- The shared disks are assigned identical drive letters in both cluster nodes.
- All IP addresses and network names in the cluster are communicating with each other and the rest of the network.
- The Cluster Service is running.
- All resources and recovery groups are online.

1 x 8 Mode on the SDS 100 Storage System

To enable clustering, the SDS 100 storage system must run in 1 x 8 mode when the two RAID controllers are connected to the system. You can verify that the backplane is in 1 x 8 mode by using the RAID controller BIOS configuration utility. Access the RAID configuration utility by pressing <Ctrl> <m> when prompted during POST. From the Management Menu, select Configure and then select View/Add Configuration. You should see the same configuration when viewing from either cluster node, particularly, the same SCSI ID numbers for the hard-disk drives, the same RAID level for the storage system, the same logical drive configuration, and so on.

If the storage system appears in 2 x 4 mode, the BIOS firmware needs to be updated on the SDS 100 storage system(s). See Appendix A, "Upgrading to a Cluster Configuration," for information about updating firmware on the SDS 100 storage system.

SCSI Controller IDs

The SDS 100 storage system of the cluster has two RAID controllers connected to the same channel. In this setup, each controller must be assigned a unique SCSI ID number. The cluster-specific firmware running on the two RAID controllers enables two controllers to reside on the same SCSI channel and operate with unique SCSI ID numbers.

If you know the version of the firmware that should be running on the RAID controller, you can verify that it is present by observing the POST message that appears during start-up that identifies the controller's firmware version. Be sure that the POST message pertains to the RAID controller connected to the shared storage subsystem.

The SCSI ID numbers on the RAID controllers can be verified using the RAID controller BIOS configuration utility. During POST, press <Ctrl> <m> to start the configuration utility. From the Management Menu, select Objects, then select Adapter, then select the appropriate adapter (if applicable), and then select Initiator ID. The SCSI IDs for the two controllers must be different from each other. The recommended settings are SCSI ID 7 for the first controller on the channel and SCSI ID 10 for the second controller on the channel.

Also, the write policy for the cluster-enabled RAID controller will be set to write-through.

Cluster Domain

On a clustered system, all systems connected to the cluster must belong to a common domain. To check that a domain has been set up properly for the cluster, start each server and client of the cluster and verify that each system can log on to the domain. To do this, go to the Control Panel, double-click on Network, and select the Identification tab. The domain name will appear in the domain field.

If the PDC does not reside in the cluster, be sure that the PDC is running before starting the systems on the cluster.

RAID Controller Driver

To verify that the PowerEdge Expandable RAID Controller driver is installed and running on the system, click the Start button, point to Settings, click Control Panel, and double-click the SCSI Adapters icon.

Click the Drivers tab and check that the PowerEdge RAID II Adapters driver shows a status of Started. Then use Windows NT Explorer to view the `winnt\system32\drivers` directory. Right-click the `pedge.sys` file, select Properties, and select the Version tab from the dialog box. Verify that the file version is 2.04 or later.

Shared Storage Subsystem Drive Letters

The shared hard-disk drives must be assigned the same drive letters in Windows NT Server Enterprise Edition running on each cluster node. The drive letters must be identical across all cluster nodes to ensure that the nodes have the same view of the file system. To check the drive letters for the shared storage subsystem(s), run the Windows NT Disk Administrator utility on one node to find the drive letters for the shared disk drives and compare the drive letters with those reported by the Disk Administrator utility on the other cluster node.

If the two systems do not see the same drive letter designation for the shared storage subsystems, the Cluster Server software was installed incorrectly. To correct this problem, uninstall the Cluster Server, reassign the drive letters, and then reinstall the Cluster Server. Refer to "Uninstalling Microsoft Cluster Server" later in this chapter for instructions.

Cluster Network Communications

For proper functioning of the cluster, the two PowerEdge systems must be able to communicate with one another. For instance, this communication includes the exchange of heartbeat messages, whereby the two servers inquire about each other's status, or "health," and acknowledge all such inquiries.

To verify network communications between the cluster nodes, open a command prompt on each cluster node. Type `ipconfig /all` at the prompt and press <Enter> to observe all known IP addresses on each local server. From each remote computer, issue the `ping` command to test the responsiveness of each IP address. Perform the same check with the cluster IP address and the IP address for each disk recovery group. Also check the cluster name and the name of each disk recovery group (if any).

Cluster Service

The Cluster Service performs most of the cluster functionality, including membership management, communication management, and fail-over management. When the Cluster Server has been properly installed, the Cluster Service is started on each node and is activated automatically in the event that one of the nodes fails or goes off-line.

To verify that the Cluster Service is running on a cluster node, click the Start button, point to Settings, and then click Control Panel. Double-click the Services icon. The Cluster Service should be indicated in the dialog box. Check to make sure that the Cluster Service is running on the second node also.

Availability of Cluster Resources

In the context of clustering, a resource is a basic unit of fail-over. Applications are made up of resources that are grouped together for the purposes of recovery. All recovery groups, and therefore their comprising resources, must be online (or in a ready state) for the cluster to function properly.

To verify that the cluster resources are online, start the Cluster Administrator on the monitoring node. Click the Start button, point to Programs, point to Administrative Tools (Common), and then click Cluster Administrator. Open a connection to the cluster and observe the running state of each recovery group. If a group has failed, one or more of its resources may be off-line.

Troubleshooting the reasons that resources might be failing is beyond the scope of this document, but examining the properties of each resource and ensuring that the specified parameters are correct is a first step in this process. In general, if a resource is off-line, it can be brought online by selecting it, right-clicking it, and choosing Bring Online from the pull-down menu. For information about troubleshooting resource failures, refer to the *Microsoft Windows NT Enterprise Edition Administrator's Guide and Release Notes*.

Uninstalling Microsoft Cluster Server

Before you can uninstall Cluster Server from a node, you must do the following:

1. **Take all resource groups off-line or move them to the other node.**
2. **Evict the node from the cluster by right-clicking the node icon in Cluster Administrator and selecting Evict Node from the menu.**
3. **Close Cluster Administrator on the node.**
4. **Stop the Cluster Service running on the node.**
5. **Uninstall Microsoft Cluster Server using the Add/Remove Programs utility in the Control Panel group.**

Removing a Node From a Cluster

WARNING: The power supplies in this computer system produce high voltages and energy hazards, which can cause bodily harm. Only trained service technicians are authorized to remove the computer covers and access any of the components inside the computer.

WARNING: This system may have more than one power supply cable. To reduce the risk of electrical shock, a trained service technician must disconnect all power supply cables before servicing the system.

When removing a node from a cluster, it is important to power down the node before removing any of the cluster cabling. Likewise, when rejoining a node to a cluster, all cables must be attached before the node is powered up.

Setting Up the Quorum Resource

A quorum resource is typically a hard-disk drive in the shared storage subsystem that serves the following two purposes in a cluster system:

- Acts as an arbiter between the two nodes to ensure that the specific data necessary for system recovery is maintained consistently across the nodes
- Logs the recovery data sent by the cluster nodes

Only one cluster node can control the quorum resource at one time, and it is that node that remains running when the two nodes are unable to communicate with each other. Once the two nodes are unable to communicate, the Cluster Service automatically shuts down the node that does not own the quorum resource.

With one of the cluster nodes down, changes to the cluster configuration database are logged to the quorum disk. The purpose of this logging is to ensure that the node that gains control of the quorum disk has access to an up-to-date version of the cluster configuration database.

Because the quorum disk plays a crucial role in the operation of the cluster, the loss of a quorum disk causes the failure of the Cluster Server. To prevent this type of failure, the quorum resource should be set up on a redundant array of hard-disk drives in the shared storage subsystem.

Using the *ftdisk* Driver

Microsoft Cluster Server does not support use of the Windows NT software-based, fault-tolerance driver **ftdisk** with any of the hard-disk drives in the shared storage subsystem. However, **ftdisk** can be used with the internal drives of the cluster nodes.

Cluster RAID Controller Functionality

The following subsections describe functional variances of standard and cluster-enabled PowerEdge Expandable RAID Controllers operating in a PowerEdge Cluster.

Rebuild Function Does Not Complete After Reboot or Power Loss

If the cluster node is rebooted or power to the node is lost while a PowerEdge Expandable RAID Controller is rebuilding a hard-disk drive, the RAID controller terminates the rebuild operation and marks the drive as failed. This also occurs if the rebuild is performed from the RAID controller basic input/output system (BIOS) configuration utility and the user exits the utility before the rebuild completes. This occurs with all versions of the PowerEdge Expandable RAID Controller firmware on both standard and cluster-enabled RAID controllers.

If the rebuild fails to complete due to a system restart, the rebuild must be reinitiated using the RAID Controller BIOS configuration utility or using the PowerEdge RAID Console program running in the Windows NT operating system.

Rebuild Rate Not Adjustable on Cluster-Enabled RAID Controller

If a hard-disk drive fails in a redundant array, you can recover the lost data by rebuilding the drive. The rate of data reconstruction is called the *rebuild rate*. The rebuild rate cannot be adjusted in a cluster-enabled RAID controller as it can in a standard RAID controller. The cluster-enabled RAID controller rebuilds drive information at a default rate.

Using the Maximize Feature in PowerEdge RAID Console

The Maximize feature of the PowerEdge RAID Console has the following functional limitations when running in the PowerEdge Cluster:

- The Maximize icon at the upper right corner of the PowerEdge RAID Console is disabled when you open the program in the PowerEdge Cluster.
- Whenever the PowerEdge RAID Console is minimized to the task bar, the right-click option to maximize the application is not available.
- Whenever the PowerEdge RAID Console is minimized to the task bar and you minimize another application, the PowerEdge RAID Console maximizes itself automatically.

Rebuild Operation in RAID Console

The following conditions apply to the way PowerEdge RAID Console handles rebuilds of hard-disk drives in a cluster environment:

- When you rebuild a failed drive, RAID Console shows the status of the drive as Rebuild but may not display the Rebuild Progress window during the rebuild process. You can verify that the rebuild is in operation by observing the activity indicator on the front panel of the SDS 100 storage system.
- During a rebuild operation, the RAID Console that issued the action reserves ownership of the channel where the failed drive is located until the rebuild is complete. Likewise if the RAID Console running on the peer server is simultaneously using that channel, it will be forced to remain with the adapter that controls the channel until the rebuild is complete. The RAID Console running on the peer server will not be able to switch to another adapter.

Chapter 4

Running Applications on a Cluster

This section provides general information about configuring and running applications software on the PowerEdge Cluster. To configure applications software, click the Start button, point to Programs, point to Administrative Tools (Common), and then click Cluster Administrator. In Cluster Administrator, open a connection to the cluster.

Before you start Cluster Administrator on either cluster node, make sure the Cluster Service has been started and a cluster has been formed. You can verify this by using the Event Viewer and looking for events logged by ClusSvc. You should see either of the following events:

Microsoft Cluster Server successfully formed a cluster on this node.

or

Microsoft Cluster Server successfully joined the cluster.

Setting Up Applications Software to Run on the Cluster

Setting up applications software to run on a cluster means establishing them as a group of cluster resources. Cluster resources are created using the New Resource wizard. The process of creating resources involves the following:

- The type of resource must be specified.
- The possible owners of the resource must be selected (the default is both nodes).
- The dependencies of the resource must be determined.
- The resource-specific parameters must be defined.

After a resource has been created, it must be brought online for access by the cluster nodes and clients.

The following subsections outline the creation and setup of three example cluster resources:

- Internet Information Server (IIS) service
- File-sharing service
- Print-spooling service

These examples are provided here to instruct you in setting up cluster resources using real applications software. Refer to the *Microsoft Windows NT Cluster Server Administrator's Guide* for more detailed information and instructions about creating cluster resources.

Internet Information Server Service

The IIS Virtual Root is one of the Microsoft Cluster Server resource types that can be used to provide fail-over capabilities for virtual root directories of IIS version 3.0 or later. The IIS Virtual Root depends on three other types of resources (disk, Internet Protocol [IP] address, and network name resources); these resources will be placed in the same recovery group.

The following example procedure describes how to set up the IIS Virtual Root service. This procedure assumes that IIS has already been installed.

1. **Start the New Group wizard by right-clicking any group or resource in the Cluster Administrator, then point to New, and then select Group from the submenu.**
2. **In the dialog box, type Web Service for the new group name.**

You may also want to select one of the cluster nodes as the preferred owner of the group.

3. **Use the New Resource wizard to create a disk resource. To start the New Resource wizard, right-click any group or resource, point to New, and then select Resource from the submenu.**

You can also move an existing disk resource from other groups by right-clicking the disk, pointing to Change Group, and then selecting Web Service.

4. **In the dialog box, type Web Disk for the new disk resource name.**
5. **Set the Resource Type in the dialog box as Physical Disk. Select both cluster nodes as possible owners and select a shared disk.**

There is no dependency for a physical disk.

NOTE: When a new resource is created, the resource group is marked off-line. This is normal and does not indicate a failure. Once the resource is created and brought online, the group is automatically brought online as well.

6. **Using the New Resource wizard, create an IP Address resource called “Web IP.”**
7. **Set the Resource Type as IP Address. Select both nodes as possible owners and then fill in an IP address and the subnet mask for your public local area network (LAN).**

There is no dependency for IP addresses.
8. **Using the New Resource wizard, create a Network Name resource called “Web NetName.”**
9. **Set the Resource Type as Network Name. Select both nodes as possible owners. Set Web IP as the dependency for Web NetName. Then type a network name that will be visible to clients (for example, website).**
10. **Use the New Resource wizard to create a IIS Virtual Root resource called “Web IIS Root.”**
11. **Set the Resource Type as IIS Virtual Root. Select both nodes as possible owners. Set Web Disk, Web IP, and Web NetName as the dependencies for Web IIS Root.**
12. **Select the WWW tab and fill in the directory and the alias in the Parameters tab.**

For example, you can configure /documents as an alias for z:\mywebdir. You should also create the same directory and place Web files there.

After bringing both the resources and the group online, users can access the IIS Virtual Root via the following URL:

http://website/documents

File Share Service

The File Share is a Cluster Server resource type that can be used to provide fail-over capabilities for file sharing. Like the IIS Virtual Root, the File Share service also depends on disk, IP address, and network name resources; these resources will be placed in the same recovery group.

The following example procedure describes how to set up the File Share service:

1. **Use the New Group wizard to create a new group called “File Share Service.”**

You may also want to select one of the cluster nodes as the preferred owner of the group.
2. **Use the New Resource wizard to create a disk resource called “File Share Disk” or move an existing shared disk resource from other groups.**
3. **Set the Resource Type in the dialog box as Physical Disk. Select both cluster nodes as possible owners and select a shared disk.**

There is no dependency for a physical disk.
4. **Use the New Resource wizard to create an IP Address resource called “File Share IP.”**
5. **Set the Resource Type as IP Address. Select both nodes as possible owners and then fill in an IP address and the subnet mask for your public LAN.**

There is no dependency for IP addresses.
6. **Use the New Resource wizard to create a Network Name resource called “File Share NetName.”**
7. **Set the Resource Type as Network Name. Select both nodes as possible owners. Set File Share IP as the dependency for File Share NetName. Then type a network name that will be visible to clients (for example, sharedfile).**
8. **Use the New Resource wizard to create a File Share resource called “XYZ Files.”**

9. **Set the Resource Type as File Share. Select both nodes as possible owners. Set File Share Disk, File Share IP, and File Share NetName as the dependencies for XYZ Files. Then type the share name and share path in the Parameters tab.**

For example, you can configure `y:\groupfiles` as share name `xyzfiles`.

NOTE: When creating a File Share resource in Microsoft Cluster Server, do not use “m” as the share name of the resource. Cluster Server rejects “m” as a File Share resource name.

After bringing both the resources and the group online, users can use Windows NT Explorer to map `xyzfiles` to a local drive.

Print Spooler Service

The Print Spooler service is a Cluster Server resource type that can be used to provide fail-over capabilities for print spooling. Like the IIS Virtual Root and the File Share service, the Print Spooler service also depends on disk, IP address, and network name resources; these resources will be placed in the same resource group.

The following example procedure describes how to set up the Print Spooler service using a HP LaserJet 5M printer. The procedure differs slightly for different printers. Make sure that Microsoft TCP/IP Printing has been installed and the printer is attached to network. Also, keep the printer's IP address and the Windows NT Server Enterprise Edition CD available.

1. **Use the New Group wizard to create a new group called “Spool Service.”**
2. **Use the New Resource wizard to create a disk resource called “Spool Disk” or move an existing shared-disk resource from other groups.**
3. **Set the Resource Type as Physical Disk. Select both cluster nodes as possible owners and then select a shared disk.**

There is no dependency for a physical disk.

4. **Use the New Resource wizard to create an IP Address resource called “Spool IP.”**

5. **Set the Resource Type as IP Address. Select both nodes as possible owners and then type an IP address and the subnet mask for your public LAN.**

There is no dependency for IP addresses.

6. **Use the New Resource wizard to create a Network Name resource called “Spool NetName.” Set the Resource Type as Network Name. Select both nodes as possible owners. Set Spool IP as the dependency for Spool NetName. Then type a network name that will be visible to clients (for example, `spoolname`).**
7. **Use the New Resource wizard to create a Print Spooler resource called “X Print.”**
8. **Set the Resource Type as Print Spooler. Select both nodes as possible owners. Set Spool Disk, Spool IP, and Spool NetName as the dependencies for “X Print.” Then type the spool folder in the Parameters tab (for example, `x:\spool`).**
9. **Bring both resources and the group online.**
10. **Install the same printer ports and printer drivers on each cluster node.**
 - a. Install the printer driver—in this example, JetAdmin for HP printers—using the installation instructions provided in your printer documentation.
 - b. After the printer driver is installed, click the Start button, point to Settings, and click Control Panel.
 - c. Double-click Printers, and then double-click Add Printer.
 - d. Select My Computer and click Next.
 - e. Click Add Port.
 - f. Highlight HP JetDirect Port and click New Port.
 - g. Either click Search to find the printer or type its IP address in the TCP/IP Address field and click Next.
 - h. Type a port name, for example `sigport`, and click Finish.
 - i. Click Close.

- j. Click Cancel to close the Add Printer wizard.
NOTE: Do not add the printer at this point. Identical printer ports must be set up on both nodes before the printer can be added.
- k. Repeat steps a through j on the other node. At step g, if the system cannot find the printer, you may need to update the HP JetAdmin's printer directory to include the printer's IP address.

11. Add the printers to the clustered spooler.

- a. On the first cluster node, click the Start button and click Run.
- b. Type \\spoolname and press <Enter>.
- c. Double-click Printers and then double-click Add Printer.
- d. Select Remote print server \\spoolname and click Next.
- e. Select sigport, then select the vendor and model, and click Next.
- f. Type a printer name, for example, sigprint, select Shared, and exit the Add Printer wizard.
- g. Click the Start button, point to Settings, and click Control Panel.
- h. Double-click Printers, and then double-click Add Printer.
- i. Select Network Printer Server and click Next.
- j. Select \\spoolname\sigprint, click OK, and then click Finish.
- k. Right-click on the sigprint icon and click Properties.
- l. Click the Scheduling tab and select Start printing after last page is spooled. Click OK to close.
- m. Repeat steps g through l on the other node.

Using the Rediscovery Application in Intel LANDesk

If the cluster node that has the system management bus (SMB) connection to the PowerEdge Scalable Disk System 100 (SDS 100) storage system fails, two actions must be taken to reestablish management from the remaining server:

1. **The SMB connection must be reestablished between the remaining cluster node and the SDS 100. That is, the cable must be physically removed from the back of the failed node and connected to the back of the remaining node.**
2. **The xover program must be run on the remaining node (now in charge of managing the SDS 100 storage system) so that the LANDesk console can rediscover its targets.**

The **xover** application is in the Server Manager installation directory (the default is `c:\smm32`). It can be run from the command line or from Windows NT Explorer. This application allows the secondary PowerEdge server to start managing the SDS 100 chassis. When the **xover** application has discovered the SDS 100, the utility acknowledges the connection with a message box. Additionally, application logs are entered in the Windows NT System Event Log by the Dell Baseboard Agent, indicating that the SDS 100 chassis has been discovered.

NOTE: xover can be run in quiet mode by specifying the -q option. In this case, the utility will not display any messages unless an error was encountered.

Running `chkdsk /f` on a Quorum Disk

The **chkdsk** command with the **/f** (“fix”) option cannot be run on a device on which an open file handle is active. The Cluster Service maintains an open handle on the quorum resource; therefore, **chkdsk /f** cannot be run on the hard-disk drive that contains the quorum resource. To run **chkdsk /f** on a quorum resource’s hard-disk drive, move the quorum resource temporarily to another drive and then run **chkdsk /f** on the drive that previously stored the quorum resource. To move the quorum resource, right-click the cluster name, select Properties, and then select the Quorum tab. Select another disk as the quorum disk and press Enter. Upon completion, move the quorum disk back to the original drive.

Tape Backup for Clustered Systems

Contact your Dell Sales representative for information about the availability of tape backup solutions and applications software for the Dell PowerEdge Cluster.

Chapter 5

Troubleshooting

This chapter provides general troubleshooting information for the Dell PowerEdge Cluster. For troubleshooting information that is specific to the Windows NT Server

Enterprise Edition operating system and the cluster software, refer to the *Microsoft Windows NT Cluster Server Administrator's Guide*.

Table 5-1. Troubleshooting

Problem	Probable Cause	Corrective Action
The cluster server nodes cannot access the Dell PowerEdge Scalable Disk System 100 (SDS 100) storage system, or the cluster software is not functioning with the storage system.	The SDS 100 storage system has not been upgraded with the cluster-specific firmware.	Ensure that the system management bus (SMB)-connected node on the cluster is running the cluster-specific firmware. Upgrade the SDS 100 firmware by powering down the cluster and then starting it up again. During start-up, the cluster-specific firmware on the node checks the version of the SDS 100 firmware. If the SDS 100 is found to be running the wrong version of firmware, the node proceeds to upgrade it automatically with the correct firmware version.
	The PowerEdge Expandable RAID Controllers have the same small computer system interface (SCSI) ID.	Set the controller in the primary node to SCSI ID 7, and set the controller in the secondary node to SCSI ID 10. Refer to Chapter 3 for instructions for setting the SCSI IDs on the nodes.
System management of the shared storage subsystem(s) is not available when the SMB-connected node fails.	The SMB connection is lost when the SMB-connected node fails and needs to be reestablished with the secondary node.	Disconnect the SMB cable from the failed node and connect it to the secondary node. Run the cluster rediscovery application in LANDesk so that the fail-over system scans the hard-disk drives on the shared storage subsystem(s) and reestablishes system management of the drives. Refer to Chapter 4 for information on running the rediscovery application.

Table 5-1. Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
The redundant array of inexpensive disks (RAID) drives in the SDS 100 storage system are not accessible by one of the cluster nodes, or the shared storage subsystem is not functioning properly with the cluster software.	The SCSI cables are loose or defective, or the cables exceed the maximum allowable length of 4 meters (m).	Check the cable connections or replace the cable with a working cable. Ensure that the length of the cable does not exceed 4 m.
	The appropriate cluster-specific PowerEdge Expandable RAID Controller driver is not running on the system.	Install the RAID controller driver that came with your system or cluster upgrade kit. Refer to Appendix A for instructions on installing the RAID controller driver.
	The RAID controllers connected to a single storage system are not configured consistently.	Ensure that the RAID configuration is identical between the RAID controllers connected to the storage system.
	If the cluster has multiple SDS 100 storage systems, the cabling between the RAID controller and the storage systems is wrong.	Be sure that the cables attached to the channel 0 connectors on the RAID controllers are connected to one storage system and the channel 1 RAID controller cables are connected to the other storage system.
	The SCSI cable between the node and the shared storage subsystem is faulty or not connected.	Attach or replace the SCSI cable between the cluster node and the shared storage subsystem.
The SDS 100 is not running in 1 x 8 mode.	A cluster-specific configuration file for the system configuration utility is missing from the system.	Install the configuration file update for the system configuration utility that came with your system or cluster upgrade kit.
	The Cluster Mode setting in the system configuration utility is incorrect for a clustered system.	Enter the system configuration utility and change the Cluster Mode field to Enabled. The Cluster Mode setting tells the cluster node to download the cluster firmware to the SDS 100 storage system, which then sets the storage system in 1 x 8 mode. Refer to Appendix A for information on updating the system configuration utility.

Table 5-1. Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
One or more of the SCSI controllers are not detected by the system.	The controllers have conflicting SCSI IDs.	Change one of the controller SCSI IDs so that the ID numbers do not conflict. The controller in the primary node should be set to SCSI ID 7, and the controller in the secondary node should be set to SCSI ID 10. Refer to Chapter 3 for instructions for setting the SCSI IDs on the nodes.
One of the nodes can access one of the shared hard-disk drives, but the second node cannot.	The drive letters assigned to the hard-disk drive differ between the nodes.	Change the drive letter designation for the shared hard-disk drive so that it is identical in all nodes.
	The SDS 100 storage system has not been upgraded with the cluster-specific firmware.	Ensure that the SMB-connected node on the cluster is running the cluster-specific firmware. Upgrade the SDS 100 firmware by powering down the cluster and then starting it up again. During start-up, the cluster-specific firmware on the node checks the version of the SDS 100 firmware. If the SDS 100 is found to be running the wrong version of firmware, the node proceeds to upgrade it automatically with the correct firmware version.
	The SCSI cable between the node and the shared storage subsystem is faulty or not connected.	Attach or replace the SCSI cable between the cluster node and the shared storage subsystem.
Server management functions are unavailable when both nodes are functional.	The SMB cable is not connected properly to the SDS 100 storage system(s).	Check the SMB connections. The primary node should be connected to the first storage system, and the second storage system (if present) should be connected to the first storage system. Refer to Chapter 2 for information about connecting the SMB cable.
Clients are dropping off of the network while the cluster is failing over.	The service provided by the recovery group becomes temporarily unavailable to clients during fail-over. Clients may lose their connection if their attempts to reconnect to the cluster are too infrequent or if they end too soon.	Reconfigure the dropped client to make longer and more frequent attempts to reconnect back to the cluster.

Table 5-1. Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
The dialogue box Snmp.exe - Entry Point Not Found appears during system start-up.	The Windows NT system errantly reports this condition if the Simple Network Management Protocol (SNMP) service was installed after Windows NT and the Service Pack was installed.	Reapply the Service Pack that came with Windows NT Enterprise Edition.
Attempts to connect to a cluster using Cluster Administrator fail.	The Cluster Service has not been started, a cluster has not been formed on the system, or the system has just been booted and services are still starting.	Verify that the Cluster Service has been started and that a cluster has been formed. Use the Event Viewer and look for the following events logged by ClusSvc: Microsoft Cluster Server successfully formed a cluster on this node. or Microsoft Cluster Server successfully joined the cluster. If these events do not appear, refer to the <i>Microsoft Cluster Server Administrator's Guide</i> for instructions on setting up the cluster on your system and starting the Cluster Service.

Appendix A

Upgrading to a Cluster Configuration

This appendix provides instructions for upgrading your noncluster system to a PowerEdge Cluster if components of the cluster hardware are already present. To properly upgrade your system to a PowerEdge Cluster, you must ensure that your existing hardware components meet the minimum configuration required for clustering and acquire the additional hardware and software clustering components as needed.

NOTES: Dell certifies only PowerEdge Cluster systems that are configured with the Dell products described in this Installation and Troubleshooting Guide (see Chapter 1 for a description of the PowerEdge Cluster components).

Dell also supports only the certification of Dell PowerEdge Cluster systems that are configured according to the instructions provided in this guide. Configurations using non-Dell products, such as server systems, rack cabinets, and storage systems, have not been approved by any safety agencies. It is the responsibility of the customer to have such systems evaluated for suitability by a certified safety agency.

After installing the necessary upgrade hardware, such as redundant array of inexpensive disk (RAID) controllers and network interface controllers (NICs), you can begin to set up and cable the system hardware. When the cluster hardware has been set up, firmware for the Dell PowerEdge 4200 systems and the PowerEdge Scalable Disk System 100 (SDS 100) storage system(s) must be updated for clustering functionality. The final phase of a PowerEdge Cluster upgrade is the installation and configuration of the Windows NT Server Enterprise Edition operating system and Cluster Server software.

Checking Your Existing Hardware

Before you can upgrade your system, you must ensure that your existing hardware meets the minimum configuration requirements for the Dell PowerEdge Cluster. See Chapter 1, “Getting Started,” for a list of the components and minimum system configuration required for the PowerEdge Cluster. Contact Dell for information on acquiring the related hardware components and customer kits that you need for the upgrade.

Chapter 1 also provides an overview of the cluster installation procedure. Refer to this chapter for the proper order of installation steps.

Adding Expansion Cards for a Cluster Upgrade

The cluster-enabled PowerEdge Expandable RAID Controller for the shared storage subsystem must be placed in Peripheral Component Interconnect (PCI) slot 7 of your PowerEdge server. If you have a second shared storage subsystem and plan to use a second cluster-enabled PowerEdge Expandable RAID Controller, install that RAID controller in PCI slot 5. Thus, two cluster-enabled RAID controllers in a cluster node occupy PCI slot 7 for the first cluster-enabled RAID controller and slot 5 for the second cluster-enabled RAID controller.

After the cluster-enabled RAID controller(s) are installed, you must disable the basic input/output system (BIOS) for these controllers. See the “Disabling a RAID Controller BIOS” in Chapter 3 for instructions.

Mounting, Cabling, and Configuring the Cluster Hardware

When you have acquired all the necessary hardware and software cluster components, you are then ready to install and connect the components into a clustered system. If you are installing the PowerEdge Cluster in a Dell Rack-Mounted Solutions enclosure, refer to Appendix B, “Stand-Alone and Rack Configurations,” for proper placement of the PowerEdge Cluster components in the rack. Instructions are also provided for installing the 3Com network switch in the rack. For further instructions for mounting Dell equipment in a Dell rack, refer to the *Dell PowerEdge Rack-Mountable Solutions Installation Guide*.

For instructions on cabling the components into a clustered system, see Chapter 2 in this guide.

You may choose to install a standard PowerEdge Expandable RAID Controller as the second RAID controller in your system instead of a cluster-enabled RAID controller. This is the configuration you will use if you plan to set up the internal hard-disk drives in the cluster node in a RAID array. In this configuration, install the standard RAID controller in PCI slot 6. Because this standard RAID controller will handle the system partition on the cluster node, do not disable this controller’s BIOS.

PCI slots 4 and 8 should be used for the cluster NIC cards. Use PCI slot 8 for the public local area network (LAN) and PCI slot 4 for the private node-to-node network.

Figure A-1 shows the back view of a cluster node with the two NICs installed in the recommended slot locations, the cluster-enabled PowerEdge Expandable RAID Controller in PCI slot 7, and a standard PowerEdge Expandable RAID Controller in PCI slot 6. Slot 5 is where you would install a second cluster-enabled RAID controller.




Figure A-1. Back View of a PowerEdge 4200 Cluster Node

Installing and Configuring the Cluster Software

Instructions on installing the Microsoft Windows NT Server, Enterprise Edition 4.0 operating system and clustering software are provided in the Microsoft documentation that accompanied the software. Refer to these documents for information about installing these software components.

Before installing Windows NT Server Enterprise Edition, be sure to have the NIC and PowerEdge Expandable RAID Controller driver diskettes handy for installation. Refer to the Microsoft documentation for information about installing the device drivers.

Before you can use the cluster software on your cluster nodes, you must first upgrade the firmware on the PowerEdge systems servers and the SDS 100 storage system(s). You must also update the system configuration utility with the cluster configuration file.

Upgrading the PowerEdge 4200 Firmware

To enable clustering on the PowerEdge system nodes, both systems must have cluster-specific BIOS firmware installed. To update the system BIOS for a PowerEdge 4200 system, perform the following steps:

1. **Insert the *Customer BIOS Update for Dell PowerEdge 4200* diskette into drive A of the PowerEdge system and restart the system.**

A message appears stating that the system is ready to update the BIOS.

2. **Type *y* to upgrade the firmware.**

While the BIOS is updating, a series of status messages appears on the screen. Another message appears when the upgrade is complete.

3. **Repeat steps 1 and 2 on the second PowerEdge system, if applicable.**

Upgrading the PowerEdge SDS 100 Storage System Firmware

The updated PowerEdge system BIOS automatically downloads the cluster-specific BIOS to the SDS 100 storage system(s) via the system management bus (SMB)

cable during the power-on self-test (POST). Upgrade the SDS 100 firmware by first ensuring that the SMB cable is connected and then powering down the cluster and then starting it up again. During start-up, the cluster-specific firmware on the SMB-connected node checks the version of the SDS 100 firmware. If the SDS 100 is found to be running the wrong version of firmware, the node proceeds to upgrade it automatically with the correct firmware version. Observe the system messages during POST to verify that the BIOS is downloading to the storage system(s).

Setting the Cluster Mode With BIOS Setup

Use the following BIOS Setup procedure to enable Cluster Mode on each node on the cluster:

1. **Start the first cluster node and press <F2> during POST.**
The Dell System PowerEdge 4200/xxx Setup screen appears.
2. **Use the right-arrow key to select the Advanced menu.**
3. **Use the down-arrow key to select the Cluster option.**
4. **Press <Spacebar> to turn on Cluster Mode.**
5. **Select the Exit menu, select Save Changes & Exit, and press <Enter>.**
6. **Restart the system.**
7. **Repeat steps 1 through 6 on the second cluster node.**

Installing and Configuring NICs

The PowerEdge Cluster requires at least two network interconnects for cluster operation—one network for the public LAN and one dedicated network for the node-to-node communications. Having two networks on the cluster enables fault-tolerance of the cluster's network communications and enables NIC replacement or upgrades without losing network connectivity.

NICs installed in the same node must reside on separate subnetworks. Therefore, the second NIC added to a cluster node must have a different network Internet Protocol (IP) address than the first NIC on the same node.

The procedure for adding and setting up a NIC in a cluster node is provided below.

NOTE: The IP addresses used are examples only and are not representative of actual addresses that should be used.

This procedure assumes that Windows NT Enterprise Edition, the current Windows NT Service Pack, and Cluster Server are installed on both cluster nodes, and the IP addresses are 143.166.110.2 for the NIC in the first node and 143.166.110.4 for the NIC in the second node. The subnet mask for both nodes is 255.255.255.0.

1. Move all cluster resources to the first cluster node.

Refer to the *Microsoft Cluster Server Administrator's Guide* for information about moving cluster resources to a specific node.

2. Power down the second cluster node and install the second NIC in that system.

Refer to the *User's Guide* for your system for instructions about installing expansion cards in your system.

3. Boot to the Windows NT Server Enterprise Edition operating system.

4. Click the Start button, point to Settings and then click Control Panel. Double-click the Network icon.

5. Install the driver for the second NIC.

6. Enter the new NIC's IP address, making sure that the network identification (ID) portion of the IP address is different from the other adapter.

For example, if the first NIC in the node had an address of 143.166.110.2 with a subnet mask of 255.255.255.0, you may enter the following IP address and subnet mask for the second NIC:

IP Address: 143.166.100.6

Subnet Mask: 255.255.255.0

7. Click OK, exit the Control Panel, and restart the node.

8. At the Windows NT desktop, click the Start button, point to Program, select Administrative Tools (Common), and then select Cluster Administrator. Click the Network tab and verify that a new resource called "New Cluster Network" has been created.

9. Move the cluster resources over to the second cluster node and repeat steps 2 through 9 on the first cluster node.

As you did with the second node, be sure to assign the new NIC with the same subnet as the second NIC of the second node (for example, 143.166.100.7).

10. In the dialog box, add a new cluster IP address resource name and assign it the same network address as the new NIC, but give the resource a unique host address.

For example, you might assign the following IP address:

IP Address: 143.166.100.8

Subnet Mask: 255.255.255.0

If the installation and IP address assignments have been performed correctly, all of the new NIC resources will appear online and will respond successfully to **ping** commands.

Appendix B

Stand-Alone and Rack Configurations

The Dell PowerEdge Cluster can be set up in a floor-standing (stand-alone) configuration or can be mounted in a Dell Rack-Mountable Solutions enclosure. Certain rules and parameters must be followed in either case to ensure that the PowerEdge Cluster is properly configured and meets safety specifications. Dell supports only PowerEdge Cluster systems that are configured according to the instructions in this appendix.

NOTES: Dell certifies only PowerEdge Cluster systems that are configured with the Dell products described in this Installation and Troubleshooting Guide (see Chapter 1 for a description of the PowerEdge Cluster components).

Dell also supports only the certification of Dell PowerEdge Cluster systems that are configured according to the instructions provided in this guide. Configurations using non-Dell products, such as server systems, rack cabinets, and storage systems, have not been approved by any safety agencies. It is the responsibility of the customer to have such systems evaluated for suitability by a certified safety agency.

Power Requirements of the PowerEdge Cluster

Refer to Chapter 2, “Cabling the Cluster Hardware,” for important information about handling the power requirements of the PowerEdge Cluster.

WARNING: Do not attempt to cable the PowerEdge Cluster to electrical power without first planning the distribution of the cluster’s electrical load across available circuits. For operation in the Americas, the PowerEdge Cluster requires two AC circuits with a minimum capacity of 20 amperes (amps) each to handle the electrical load of the system. Do not allow the electrical load of the system to exceed 16 amps on either circuit. For operation in Europe, the PowerEdge Cluster requires two circuits rated in excess of the combined load of the attached systems. Please refer to the ratings marked on the back of each cluster component when determining the total system’s electrical load.

WARNING: Although each component of the PowerEdge Cluster meets leakage current safety requirements, the total leakage current may exceed the maximum that is permitted when the components are used together. To meet safety requirements in the Americas, you must use a Type B plug and socket connection for the cluster power to enable the appropriate level of ground protection. In Europe, you must use one (or two) power distribution units (PDUs) or two Type B plug and socket connections wired and installed by a qualified electrician in accordance with the local wiring regulations.

Supported Stand-Alone Configurations

Figures B-1 and B-2 show the stand-alone configurations of the PowerEdge Cluster that Dell supports. These configurations are specified to provide a safe operating environment for the cluster components.

As evident in the figures, two general rules govern the stand-alone configurations:

- The heaviest cluster component must be at the bottom of the stack and the lightest component at the top.
- No more than a single PowerEdge Scalable Disk System 100 (SDS 100) storage system or a single network switch can be stacked on top of a PowerEdge server. If stacked alone, the storage system(s) and network switch may be stacked one on top of the other with the network switch on top, as shown in Figure B-2.

NOTE: Placement of the monitor, keyboard, or mouse on top of the PowerEdge systems or the SDS 100 storage system(s) is not supported by Dell. Also, Dell does not support more than one network switch stacked in any stand-alone configuration of the cluster components.

Rack Safety Notices

Before you begin installing the PowerEdge Cluster components in your rack, carefully read the safety precautions and installation restrictions in the following subsections.

Kit Installation Restrictions

WARNING: Dell's server systems are certified as components for use in Dell's rack cabinet using the Dell Customer Rack Kit. The final installation of Dell servers and rack kits in any other brand of rack cabinet has not been approved by any safety agencies. It is the customer's responsibility to have the final combination of Dell servers and rack kits for use in other brands of rack cabinets evaluated for suitability by a certified safety agency.

This rack kit is intended to be installed in a Dell Rack-Mountable Solutions enclosure by trained service technicians. If you install the kit in any other rack, be sure that the rack meets the specifications of the Dell rack.


Rack Stabilizer Feet

WARNING: Installing a PowerEdge system in a Dell rack without the front and side stabilizer feet installed could cause the rack to tip over, resulting in bodily injury. Therefore, always install the stabilizer feet before installing components in the rack.



Refer to the *Dell PowerEdge Rack-Mountable Solutions Installation Guide* provided with the rack for instructions on installing the stabilizer feet.

WARNING: After installing computers in a rack, never pull more than one computer out of the rack on its slides at one time. The weight of more than one computer extended on slides could cause the rack to tip over and cause bodily injury.

PowerEdge SDS 100
storage system




network switch





Configuration 1

SDS 100 storage system



network switch



Configuration 2

Figure B-1. Supported Stand-Alone Configurations With One SDS 100 Storage System




Figure B-2. Supported Stand-Alone Configurations With Two SDS 100 Storage Systems

Supported Rack Configuration

Dell supports one configuration of the Dell PowerEdge Cluster mounted in a Dell rack. The following list shows how the cluster components must be placed in this configuration, from the lowest rack position to the highest:

- Uninterruptible power supply (UPS) (lowest rack position)
- First PowerEdge system
- Keyboard tray and network switch mounted behind the keyboard tray

- Apex Outlook switch box
- Second PowerEdge system
- First PowerEdge SDS 100 storage system
- Optional second SDS 100 storage system (highest rack position)

Figure B-3 illustrates the Dell-supported rack configuration. For instructions on installing the individual components of the PowerEdge Cluster in a Dell rack, refer to the *Dell PowerEdge Rack-Mountable Solutions Installation Guide*. Instructions on installing the network switch in a rack are provided in the next section.




Figure B-3. Supported Rack Configuration

Rack-Mounting the Network Switch

For the 3Com Superstack II Switch 3000 TX to be accessible to the network interface controller (NIC) connectors on each cluster node, the switch must be placed behind the keyboard tray with the front of the switch facing toward the back of the rack.

Use the following procedure to install the network switch in the rack.

CAUTION: Do not connect cables to the network switch prior to installing the switch in the rack.

1. If present, remove all self-adhesive pads from the underside of the network switch.
2. At the back of the rack along one of the vertical rails, locate the first square hole right below the mounting rail for the Apex Outlook switch box (which connects the mouse, keyboard, and monitor), and install a nut cage in the hole. Moving downward, skip the second and third holes and install a nut cage in the fourth hole.
3. Attach two cage nuts on the back vertical rail on the other side of the rack directly opposite of the two cage nuts you just installed.
4. With the front of the network switch facing you, position the mounting bracket over the mounting holes on one side of the switch, as shown in Figure B-4.




Figure B-4. Attaching the Rack-Mounting Hardware on the Network Switch

5. Insert three of the screws included with the mounting hardware and tighten them securely.
6. Insert the mounting bracket on the opposite side of the unit.
7. Position the network switch inside the rack behind the keyboard tray with the front of the switch facing toward the back of the rack (see Figure B-3).
8. Align the holes on the mounting hardware with the cage nuts that were installed earlier and secure them with #10 flat washers and #10-32 screws.

Refer to Chapter 2, “Cabling the Cluster Hardware,” for instructions on cabling the network switch.

Appendix C

Cluster Data Sheet

The data sheet on the following page is provided for the system installer to tear out and use to record pertinent information about the Dell PowerEdge Cluster. Have this form available when calling Dell Technical Assistance.

Dell® PowerEdge® Cluster Installer Data Card and Checklist

Instructions: Before installing the Microsoft® Windows NT® Enterprise Edition operating system with clustering, use this checklist to gather information and ensure the preparation required for a successful installation. Ensure that all equipment is present and properly cabled and that you know how to install Windows NT Server Enterprise Edition.

Cluster order number: _____

Software and Firmware

- Dell PowerEdge Expandable RAID Controller cluster firmware revision: _____
- The Intel Pro100B driver revision is 2.22 or later.
- Cluster system configuration utility version: _____
- Windows NT Server Enterprise Edition revision: _____
- xover.exe** revision: _____

Pre-Installation Settings

- PowerEdge Expandable RAID Controller initiator IDs:
Node 1: _____ and Node 2: _____
- The cluster RAID controller has been configured for write-through operation.
- The RAID level and logical drives are configured and initialized.
- Cluster Mode for the Dell PowerEdge 4200 Cluster has been turned on in the system BIOS.

Windows NT Server Settings

- The RAID controller drivers are installed.
- A 1024-MB partition has been created for the Windows NT Server Enterprise Edition system drive.
- License type: _____ Number: _____
- Network name for this computer: _____
- Domain type has been chosen (Primary Domain Controller, Backup Domain Controller, or stand-alone).
- Administrator user name: _____
- Administrator password: _____
- Network participation has been wired to network.
- Microsoft Internet Information Server (IIS) has been installed.
- All network adapters have been found and accepted.
- Network protocol is TCP/IP only. Uncheck any others.
- SNMP service has been added.
- DHCP server is not selected.

- TCP/IP address for each NIC:
Node 1, NIC 1 _____ NIC 2 _____
Node 2, NIC 1 _____ NIC 2 _____
- The subnet masks for NIC 1 and NIC 2 are different. Subnet masks for the NIC 1s should match. Subnet masks for the NIC 2s should match.
Node 1, NIC 1 _____ NIC 2 _____
Node 2, NIC 1 _____ NIC 2 _____
- Domain name: _____
- Gopher service is not enabled for IIS.
- Drive letters for the SDS 100s:
No. 1 _____ No. 2 _____ No. 3 _____ No. 4 _____
- The format of the SDS 100 logical drive is NTFS.
- The NIC driver has been updated.
- The recommended paging file maximum size is twice the system's RAM capacity if the RAM capacity is 256 MB or less, or the paging file size has been set at an amount greater than the system RAM, up to the amount of free space on the hard-disk drive.
- The recommended registry file size is 64 MB.
- The OEM NIC driver remains installed after the Service Pack installation.
- On Node 2, the drive letter for the SDS 100 is the same as on Node 1.

Microsoft Cluster Service Installation

- Cluster name: _____
- Domain name: _____
- Administrator's user name: _____
- Administrator's password: _____
- Name of adapter 1 is Public.
- Name of adapter 2 is Private.
- Cluster IP address: _____
- Cluster subnet mask (same as Public): _____

© Dell Computer Corporation 1997, Rev.1.1

Appendix D

PowerEdge Cluster Configuration Matrix

The Cluster Configuration Matrix form, which is shown on the following page, is attached to the back of each cluster node and is used by the system installer to record important information about the hardware on each cluster component. Keep these completed forms attached to their respective cluster components. Have these forms handy any time you call Dell for technical support.

The form provides fields for the following information:

- Form completion date
- Unique cluster identification (ID) number
- Service tag numbers for each cluster component

- List of each cluster node's Peripheral Component Interconnect (PCI) slots and the adapters installed in each
- Usage description for each installed adapter
- PowerEdge Scalable Disk System 100 (SDS 100) storage system service tags associated with each PowerEdge Expandable redundant array of inexpensive disks (RAID) Controller channel

NOTE: Currently, only two PowerEdge SDS 100 storage systems are supported on the PowerEdge Cluster. Future enhancements will provide support for up to four SDS 100 storage systems.

PowerEdge Cluster Configuration Matrix

Unique Cluster ID: _____

System	Service Tag #	Node Number	RAID Controller ID
PE 4200		Node 1	7
PE 4200		Node 2	10
SDS 100			
SDS 100			
SDS 100			
SDS 100			

Slot #	PCI/EISA Slot #	Adapter	Usage	Attachment Instruction	
1	EISA slot 1				
2	EISA slot 2				
3	EISA slot 3				
4	PCI slot 4 (Secondary bus)	Private NIC	Private Network		
5	PCI slot 5 (Secondary bus)	Second Cluster PERC (if applicable)	Shared Drives	Channel 0	Channel 1
6	PCI slot 6 (Primary bus)	Standard PERC (optional)	Server Drives	Channel 0	Channel 1
7	PCI slot 7 (Primary bus)	First Cluster PERC (required)	Shared Drives	Channel 0	Channel 1
8	PCI slot 8 (Primary bus)	Public NIC	Public Network		

Figure D-1. PowerEdge Cluster Configuration Matrix

Appendix E

Regulatory Compliance

Regulatory Standards

Dell PowerEdge Cluster has been tested and certified to the following standards:

- IEC950:1991 + A1:1992 + A2:1993 + A3:1995
- EN60950:1992 + A1:1993 + A2:1993 + A3:1995
- EMKO-TSE (74-SEC) 207/94
- Miti Ordinance No. 85
- UL1950, 3rd Edition
- C22-2 No. 950, 3rd Edition

CE Notice

Marking by the symbol **CE** indicates compliance of this Dell system to the Safety and EMC (Electromagnetic Compatibility) directives of the European Community (89/336/EEC and 73/23/EEC). Such marking is indicative that this Dell system meets or exceeds the following technical standards.

Safety Standard

- EN60950:1992 + Amd.1:1993 + Amd.2:1993 — “Safety of Information Technology Equipment including Electrical Business Equipment.”

EMC Standards

- EN 55022 — “Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment.”

NOTE: EN 55022 emissions requirements provide for two classifications—Class A and Class B. If any one of the registration labels (on the bottom or back panel of your computer, on card-mounting brackets, or on the cards themselves) carries an FCC Class A rating, the following warning applies to your system.

WARNING: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

- EN 50082-1 — “Electromagnetic compatibility—Generic immunity standard Part 1: Residential, commercial, and light industry.”
- IEC 801-2 — “Electromagnetic compatibility for industrial-process measurement and control equipment Part 2: Electrostatic discharge requirements.” — Severity level 3.
- IEC 801-3 — “Electromagnetic compatibility for industrial-process measurement and control equipment Part 3: Radiated electromagnetic field requirements.” — Severity level 2.
- IEC 801-4 — “Electromagnetic compatibility for industrial-process measurement and control equipment Part 4: Electrical fast transient/burst requirements.” — Severity level 2.

A “Declaration of Conformity” in accordance with the preceding standards has been made and is on file at Dell Products Europe BV, Limerick, Ireland.

Appendix F

Safety Information for Technicians

Before you perform any procedures on the PowerEdge Cluster equipment, read the following warnings for your personal safety and to prevent damage to the system from electrostatic discharge (ESD). Refer to the appropriate system documentation before servicing any system.

WARNING: The components of this cluster system may have more than one power supply cable. To reduce the risk of electrical shock, a trained service technician must disconnect all power supply cables before servicing any system components.

WARNING FOR YOUR PERSONAL SAFETY AND PROTECTION OF THE EQUIPMENT:

Before you start to work on a system component, perform the following steps in the sequence listed:

1. Turn off the system component.
2. Disconnect the system component from its power source(s).
3. Disconnect any communications cables.
4. Wear a wrist grounding strap, and clip it to an unpainted metal surface, such as a part of the back panel, on the chassis.
5. If a wrist grounding strap is not available, touch the fan guard or some other unpainted metal surface on the back of the chassis to discharge any static charge from your body.

Appendix G

Warranties and Return Policy

Limited Three-Year Warranty (U.S. and Canada Only)

Dell Computer Corporation (“Dell”) manufactures its hardware products from parts and components that are new or equivalent to new in accordance with industry-standard practices. Dell warrants that the hardware products it manufactures will be free from defects in materials and workmanship. The warranty term is three years beginning on the date of invoice, as described in the following text.

Damage due to shipping the products to you is covered under this warranty. Otherwise, this warranty does not cover damage due to external causes, including accident, abuse, misuse, problems with electrical power, servicing not authorized by Dell, usage not in accordance with product instructions, failure to perform required preventive maintenance, and problems caused by use of parts and components not supplied by Dell.

This warranty does not cover any items that are in one or more of the following categories: software; external devices (except as specifically noted); accessories or parts added to a Dell system after the system is shipped from Dell; accessories or parts added to a Dell system through Dell’s system integration department; accessories or parts that are not installed in the Dell factory; or DellWare® products. Monitors, keyboards, and mice that are Dell-branded or that are included on Dell’s standard price list are covered under this warranty; all other monitors, keyboards, and mice (including those sold through the DellWare program) are not covered. Batteries for portable computers are covered only during the initial one-year period of this warranty.

Coverage During Year One

During the one-year period beginning on the invoice date, Dell will repair or replace products covered under this limited warranty that are returned to Dell’s facility. To request warranty service, you must call Dell’s Customer Technical Support within the warranty period. Refer to the chapter titled “Getting Help” in your system *Installation and Troubleshooting Guide* to find the appropriate telephone number for obtaining customer assistance. If warranty service is required, Dell will issue a Return Material Authorization Number. You must ship the products back to Dell in their original or equivalent packaging, prepay shipping charges, and insure the shipment or accept the risk of loss or damage during shipment. Dell will ship the repaired or replacement products to you freight prepaid if you use an address in the continental U.S. or Canada, where applicable. Shipments to other locations will be made freight collect.

NOTE: Before you ship the product(s) to Dell, back up the data on the hard-disk drive(s) and any other storage device(s) in the product(s). Remove any removable media, such as diskettes, CDs, or PC Cards. Dell does not accept liability for lost data or software.

Dell owns all parts removed from repaired products. Dell uses new and reconditioned parts made by various manufacturers in performing warranty repairs and building replacement products. If Dell repairs or replaces a product, its warranty term is not extended.

Coverage During Years Two and Three

During the second and third years of this limited warranty, Dell will provide, on an exchange basis and subject to Dell's Exchange Policy in effect on the date of the exchange, replacement parts for the Dell hardware product(s) covered under this limited warranty when a part requires replacement. You must report each instance of hardware failure to Dell's Customer Technical Support in advance to obtain Dell's concurrence that a part should be replaced and to have Dell ship the replacement part. Dell will ship parts and prepay the shipping costs if you use an address in the continental U.S. or Canada, where applicable. Shipments to other locations will be made freight collect. Dell will include a prepaid shipping container with each replacement part for your use in returning the replaced part to Dell. Replacement parts are new or reconditioned. Dell may provide replacement parts made by various manufacturers when supplying parts to you. The warranty term for a replacement part is the remainder of the limited warranty term.

You will pay Dell for replacement parts if the replaced part is not returned to Dell. The process for returning replaced parts, and your obligation to pay for replacement parts if you do not return the replaced parts to Dell, will be in accordance with Dell's Exchange Policy in effect on the date of the exchange.

You accept full responsibility for your software and data. Dell is not required to advise or remind you of appropriate backup and other procedures.

General

DELL MAKES NO EXPRESS WARRANTIES BEYOND THOSE STATED IN THIS WARRANTY STATEMENT. DELL DISCLAIMS ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. SOME STATES (OR JURISDICTIONS) DO NOT ALLOW LIMITATIONS ON IMPLIED WARRANTIES, SO THIS LIMITATION MAY NOT APPLY TO YOU.

DELL'S RESPONSIBILITY FOR MALFUNCTIONS AND DEFECTS IN HARDWARE IS LIMITED TO REPAIR AND REPLACEMENT AS SET FORTH IN THIS WARRANTY STATEMENT. THESE WARRANTIES GIVE YOU SPECIFIC LEGAL RIGHTS, AND

YOU MAY ALSO HAVE OTHER RIGHTS, WHICH VARY FROM STATE TO STATE (OR JURISDICTION).

DELL DOES NOT ACCEPT LIABILITY BEYOND THE REMEDIES SET FORTH IN THIS WARRANTY STATEMENT OR LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, INCLUDING WITHOUT LIMITATION ANY LIABILITY FOR PRODUCTS NOT BEING AVAILABLE FOR USE OR FOR LOST DATA OR SOFTWARE.

SOME STATES (OR JURISDICTIONS) DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE PRECEDING EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

These provisions apply to Dell's limited three-year warranty only. For provisions of any service contract covering your system, refer to the separate service contract that you will receive.

If Dell elects to exchange a system or component, the exchange will be made in accordance with Dell's Exchange Policy in effect on the date of the exchange.

NOTE: If you chose one of the available warranty and service options in place of the standard limited three-year warranty described in the preceding text, the option you chose will be listed on your invoice.

"Total Satisfaction" Return Policy (U.S. and Canada Only)

If you are an end-user customer who bought products directly from a Dell company, you may return them to Dell up to 30 days from the date of invoice for a refund of the product purchase price if already paid. This refund will not include any shipping and handling charges shown on your invoice. If you are an organization who bought the products from Dell under a written agreement with Dell, there may be different terms for the return of products under this policy, based on your agreement with Dell. To return products, you must call Dell Customer Service (at the telephone number shown in the chapter titled "Getting Help" in your *PowerEdge 4200 Systems Installation and Troubleshooting Guide*) to receive a Credit Return Authorization Number. You must ship the products to Dell in their original packaging, prepay shipping charges, and insure the shipment or accept the

risk of loss or damage during shipment. You may return software for refund or credit only if the sealed package containing the diskette(s) or CD(s) is unopened. Returned products must be in as-new condition, and all of the manuals, diskette(s), CD(s), power cables, and other items included with a product must be returned with it.

This “Total Satisfaction” Return Policy does not apply to DellWare products, which may be returned under DellWare’s then-current return policy.

Index

A

- Apex switch cabling, 2-7
- applications
 - setting up to run on a cluster, 4-1
- audience level, ix

B

- basic input/output system. *See* BIOS
- BIOS
 - disabling on a RAID controller, 3-2

C

- cabling
 - disconnecting SCSI cables, 2-7
 - mouse, keyboard, and monitor, 2-7
 - NIC, 2-5
 - SDS 100 storage systems, 2-1, 2-3, 2-4
 - SMB, 2-5
- cautions, x
- chapter summaries, ix
- chkdsk /f
 - running on a quorum disk, 4-5
- cluster
 - cabling, 2-1
 - checking the functionality, 3-5
 - components, 1-1

- cluster (*continued*)
 - configuring the software, 3-1
 - running applications on, 4-1
 - troubleshooting, 5-1
 - verifying network communications, 3-6
- cluster layout, 1-2
- cluster node
 - adding peripherals, 1-4
 - back view, 1-5
 - changing the IP address, 3-4
 - removing from cluster, 3-7
- cluster resources
 - verifying availability, 3-7
- cluster service
 - verifying operation, 3-7
- cluster software
 - high-level configuration, 3-3
 - low-level configuration, 3-1
- conventions used in text, xi

D

- domain
 - choosing for the cluster, 3-3
 - verifying operation, 3-6
- drive letters
 - assigning to shared drives, 3-4

E

electrostatic discharge. *See* ESD

ESD

- about, vi
- preventing, vi, F-1

expansion cards

- placement on PCI bus, 1-4, A-1

F

File Share service, 4-2

ftdisk driver, 3-8

G

getting started, 1-1

I

IIS, 4-1

installation overview, 1-3

internal hard-disk drive

- setting the RAID level, 3-2

Internet Information Server Service. *See* IIS

IP address

- changing for a cluster node, 3-4
- requirements for cluster, 3-3

K

keyboard cabling, 2-7

M

maximize feature in RAID Console, 3-8

Microsoft Cluster Server

- uninstalling, 3-7

monitor cabling, 2-7

mouse cabling, 2-7

N

network

- configuring separate networks, 3-3

network communications

- verifying, 3-6

network interface controller. *See* NIC

network switch

- attaching rack-mounting hardware, B-6
- rack-mounting, B-6

NIC

- cabling, 2-5
- installing, A-3
- location on PCI bus, 1-4, A-1
- updating the driver, 3-5

notational conventions, x

notes, x

P

paging file size, 3-5

PCI slots

- expansion card placement, 1-4, A-1

peripherals

- adding expansion cards for clustering, 1-4

power requirements, B-1

- PowerEdge Cluster
 - checking the functionality, 3-5
 - components, 1-1
 - getting started, 1-1
 - installation overview, 1-3
 - layout, 1-2
 - minimum system requirements, 1-2

- PowerEdge RAID Console
 - rebuild operation, 3-9
 - using the maximize feature, 3-8

- PowerEdge Scalable Disk System 100. *See* SDS 100 storage system

- Print Spooler service, 4-3

Q

- quorum resource
 - setting up, 3-8

R

- rack configuration
 - Dell-supported, B-5
 - rack certification, B-2, B-6
 - safety notices, B-2
 - stability warnings, B-2, B-6

- RAID controller
 - disabling the BIOS, 3-2
 - driver, 3-4
 - functionality, 3-8
 - location on PCI bus, 1-4, A-1
 - setting SCSI IDs, 3-2
 - verifying the driver, 3-6

- RAID level
 - setting for internal hard-disk drives, 3-2
 - setting for shared storage subsystem, 3-2

- rebuild
 - incomplete, 3-8
 - operation in RAID Console, 3-9
 - rate, 3-8
- rediscovery application, 4-4
- registry size, 3-5
- return policy, G-2

S

- safety
 - information for technicians, F-1
 - notices, B-2
 - preventing electrostatic discharge, vi

- SCSI
 - disconnecting cables, 2-7

- SCSI ID
 - setting, 3-2
 - verifying, 3-6

- SDS 100 storage system
 - cabling, 2-1, 2-3, 2-4
 - upgrading the firmware, A-3
 - verifying operation, 3-5

- shared storage subsystem
 - assigning drive letters, 3-4
 - setting the RAID level, 3-2
 - verifying drive letters, 3-6

- small computer system interface. *See* SCSI

- SMB
 - cabling, 2-5

- stand-alone configurations
 - containing one SDS 100 storage system, B-3
 - containing two SDS 100 storage systems, B-4
 - Dell-supported, B-2

- system configuration utility
 - updating for clustering, A-3

- system management bus. *See* SMB

- system requirements, 1-2

T

tape backup for clustered systems, 4-5

troubleshooting

- cluster mode failure, 5-2
- connecting to a cluster, 5-4
- network communications, 5-3
- SCSI controllers, 5-3
- shared storage subsystem, 5-1, 5-2, 5-3
- SNMP service, 5-4
- system management bus, 5-1, 5-3

typographical conventions, xi

U

upgrading

- checking existing hardware, A-1
- existing system to a cluster, A-1
- installing hardware, A-2
- installing software, A-3
- PowerEdge 4200 firmware, A-3
- SDS 100 storage system firmware, A-3
- system configuration utility, A-3

V

verifying

- 1 x 8 mode on shared storage subsystem, 3-5
- cluster domain, 3-6
- cluster resource availability, 3-7
- cluster service operation, 3-7
- network communications, 3-6
- RAID controller driver, 3-6
- SCSI controller IDs, 3-6
- shared storage subsystem drive letters, 3-6

W

warnings, x

warranty information, G-1

working inside the computer
safety precautions, vi, F-1



www.dell.com

Printed in the U.S.A.

P/N 17088

