





Dell™ PowerEdge™ Expandable RAID Controller 5/i and 5/E User's Guide

Model UCP-50 and UCP-51

Dell™ PowerEdge™ Expandable
RAID Controller 5/i and 5/E
User's Guide

Notes, Notices, and Cautions

-  **NOTE:** A NOTE indicates important information that helps you make better use of your system.
-  **NOTICE:** A NOTICE warns against either potential damage to hardware or of data and tells you how to avoid the problem.
-  **CAUTION:** A CAUTION indicates a potential for property damage, personal injury, or death.

-  **NOTE:** See the *Product Information Guide* that came with your system for complete information about U.S. Terms and Conditions of Sale, Limited Warranties, and Returns, Export Regulations, Software License Agreement, Safety, Environmental and Ergonomic Instructions, Regulatory Notices, and Recycling Information.

Information in this document is subject to change without notice.
© 2006–2007 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, *PowerVault*, *Dell Precision*, and *OpenManage* are trademarks of Dell Inc. *MegaRAID* is a registered trademark of LSI Logic Corporation. *Microsoft*, *MS-DOS*, *Windows Server*, and *Windows* are registered trademarks, and *Windows Vista* is a trademark of Microsoft Corporation. *Intel* is a registered trademark of Intel Corporation. *Novell* and *SUSE* are registered trademarks of Novell, Inc. in the United States and other countries. *Red Hat* and *Red Hat Enterprise Linux* are registered trademarks of Red Hat, Inc.

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

Model UCP-50 and UCP-51

January 2007 P/N GD740 Rev. A03

Contents

	CAUTION: Safety Instructions	9
	SAFETY: General	9
	SAFETY: When Working Inside Your System	9
	Protecting Against Electrostatic Discharge	10
	SAFETY: Battery Disposal	10
1	Overview	11
	Scope of the User's Guide	11
	PERC 5 Controller Descriptions	11
	PCI Architecture	12
	Operating System Support	12
	RAID Description	12
	Summary of RAID Levels	12
	RAID Terminology	13
	Disk Striping	13
	Disk Mirroring	13
	Parity	14
2	PERC 5 Features	15
	PERC 5 Controller Features	15
	Compatibility With Virtual Disks Created on Existing PERC 5 Controllers	16
	SMART Technology	17
	Background Initialization	17
	LED Operation	17
	Disk Roaming	18
	Disk Migration	18
	Alarm Alert in Case of Physical Disk Failures	18

Battery Management	19
Introduction to Write Cache Policy	19
Write-Back versus Write-Through	19
How Firmware Manages Cache	19
Conditions Under Which Write-Back is Employed	20
Conditions Under Which Write-Through is Employed	20
Conditions Under Which Forced Write-Back With No Battery is Employed	20
RAID Configuration Information	21
Fault Tolerance Features	21
Physical Disk Hot Swapping	22
Patrol Read	22
Patrol Read Behavior	22
Configuration	23
Patrol Read Modes	23
Behavior Details	23
Blocked Operations	23
3 Hardware Installation and Configuration	25
Installing the PERC 5/E Adapter	25
Installing the Transportable Battery Backup Unit (TBBU) for PERC 5/E	27
Installing the DIMM on a PERC 5/E Adapter	28
Transferring a TBBU Between Controllers	30
Removing the PERC 5/E Adapter	30
Removing the DIMM and Battery from a PERC 5/E Adapter	31
Installing the PERC 5/i Adapter	33
Removing the PERC 5/i Adapter	35
Removing the DIMM from a PERC 5/i Card	37
Disconnecting the BBU from a PERC 5/i Adapter or a PERC 5/i	38

4	Driver Installation	39
	Creating a Driver Diskette	39
	Installing the Driver During a Microsoft Operating System Installation	40
	Installing a Windows Server 2003, Windows Vista, Windows 2000, or XP Driver for a New RAID Controller	40
	Updating an Existing Windows 2000, Windows Server 2003, Windows XP, or Windows Vista Driver	41
	Installing Windows Vista for a New RAID Controller.	42
	Installing the Red Hat Enterprise Linux Driver	43
	Creating a Driver Diskette	43
	Installing the Driver	44
	Installing the Driver Using an Update RPM.	44
	Installing the SUSE Linux Enterprise Server (Version 9 or 10) Driver.	45
	Installing the Driver Using an Update RPM.	45
	Installing SUSE Linux Enterprise Server Using the Driver Update Diskette	46
5	RAID Configuration and Management	49
	Dell OpenManage Storage Management	49
	SAS RAID Storage Manager.	49
	BIOS Configuration Utility	49
	Entering the BIOS Configuration Utility	50
	Starting the BIOS Configuration Utility	50
	Exiting the Configuration Utility.	50
	Menu Navigation Controls.	50
	RAID Configuration Functions	52


BIOS Configuration Utility Menu Options	53
Virtual Disk Management (VD Mgmt)	54
Physical Disk Management (PD Mgmt)	56
Physical Disk Actions	57
Controller Management (Ctrl Mgmt)	57
Controller Management Disk Actions	57
Foreign Configuration View	58
Setting Up Virtual Disks	58
Creating Virtual Disks	60
Initializing Virtual Disks	61
Importing or Clearing Foreign Configurations Using the VD Mgmt Menu	62
Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen	62
Setting LED Blinking	64
Managing Dedicated Hot Spares	64
Creating Global Hot Spares	65
Removing Global or Dedicated Hot Spares	65
Enabling the Alarm to Alert in Case of Physical Disk Failures	66
Checking Data Consistency	66
Stopping Background Initialization	67
Performing a Manual Rebuild of an Individual Physical Disk	67
Deleting Virtual Disks	67
Deleting Disk Groups	68
Upgrading Firmware	68
Enabling Boot Support	68
Enabling BIOS Stop on Error	69
Restoring Factory Default Settings	69
6 Troubleshooting	71
Virtual Disks Degraded	71
Memory Errors	71
General Problems	72
Physical Disk Related Issues	72
Physical Disk Failures and Rebuilds	73
SMART Error	74

	PERC 5 Post Error Messages	75
	Red Hat Enterprise Linux Operating System Errors	78
	LED Behavior Patterns	79
	Audible Alarm Warnings	79
7	Appendix: Regulatory Notices	81
	Regulatory Notices	81
	FCC Notices (U.S. Only)	82
	FCC, Class A	82
	FCC, Class B	82
	Industry Canada (Canada Only)	83
	Industry Canada, Class A	83
	Industry Canada, Class B	83
	CE Notice (European Union)	84
	CE Notice (European Union)	84
	MIC Notice (Republic of Korea Only)	92
	VCCI (Japan Only)	93
	CNCA (China Only)	94
	CNCA, Class A	94
	información NOM (únicamente para México)	95
	Corporate Contact Details (Taiwan Only)	95
	Glossary	99
	Index	109

CAUTION: Safety Instructions

Use the following safety guidelines to help ensure your own personal safety and to help protect your system and working environment from potential damage.

 **CAUTION:** 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. See "SAFETY: Battery Disposal" on page 10.

 **NOTE:** See the safety regulations and warnings stated in the documentation that ships with your PowerEdge™ system/Dell Precision™ workstation.


SAFETY: General

- Observe and follow service markings. Do not service any product except as explained in your user documentation. Opening or removing covers that are marked with the triangular symbol with a lightning bolt may expose you to electrical shock. Components inside these compartments should be serviced only by a trained service technician.
- If any of the following conditions occur, unplug the product from the electrical outlet and replace the part or contact your trained service provider:
 - The power cable, extension cable, or plug is damaged.
 - An object has fallen into the product.
 - The product has been exposed to water.
 - The product has been dropped or damaged.
 - The product does not operate correctly when you follow the operating instructions.
- Use the product only with approved equipment.
- Operate the product only from the type of external power source indicated on the electrical ratings label. If you are not sure of the type of power source required, consult your service provider or local power company.
- Handle batteries carefully. Do not disassemble, crush, puncture, short external contacts, dispose of in fire or water, or expose batteries to temperatures higher than 60 degrees Celsius (140 degrees Fahrenheit). Do not attempt to open or service batteries; replace batteries only with batteries designated for the product.

SAFETY: When Working Inside Your System

Before you remove the system covers, perform the following steps in the sequence indicated.

 **CAUTION:** Except as expressly otherwise instructed in Dell documentation, only trained service technicians are authorized to remove the system cover and access any of the components inside the system.

 **NOTICE:** 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.

- 1 Turn off the system and any devices.
- 2 Ground yourself by touching an unpainted metal surface on the chassis before touching anything inside the system.
- 3 While you work, periodically touch an unpainted metal surface on the chassis to dissipate any static electricity that might harm internal components.
- 4 Disconnect your system and devices from their power sources. To reduce the potential of personal injury or shock, disconnect any telecommunication lines from the system.

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

- 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 that both connectors are correctly oriented and aligned.
- Handle components and cards with care. Do not 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.

Protecting Against Electrostatic Discharge

Electrostatic discharge (ESD) events can harm electronic components inside your computer. Under certain conditions, ESD may build up on your body or an object, such as a peripheral, and then discharge into another object, such as your computer. To prevent ESD damage, you should discharge static electricity from your body before you interact with any of your computer's internal electronic components, such as a memory module. You can protect against ESD by touching a metal grounded object (such as an unpainted metal surface on your computer's I/O panel) before you interact with anything electronic. When connecting a peripheral (including handheld digital assistants) to your computer, you should always ground both yourself and the peripheral before connecting it to the computer. In addition, as you work inside the computer, periodically touch an I/O connector to remove any static charge your body may have accumulated.

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

- When unpacking a static-sensitive component from its shipping carton, do not remove the component from the antistatic packing material until you are ready to install the component. Just before unwrapping the antistatic package, 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 electrostatic sensitive components in a static-safe area. If possible, use antistatic floor pads and work bench pads.

SAFETY: Battery Disposal



Your system may use a nickel-metal hydride (NiMH), lithium coin-cell, and/or a lithium-ion battery. The NiMH, lithium coin-cell, and lithium-ion batteries are long-life batteries, and it is very possible that you will never need to replace them. However, should you need to replace them, refer to the instructions included in the section "RAID Configuration and Management" on page 49.

Do not dispose of the battery along with household waste. Contact your local waste disposal agency for the address of the nearest battery deposit site.



NOTE: Your system may also include circuit cards or other components that contain batteries. These batteries must also be disposed of in a battery deposit site. For information about such batteries, refer to the documentation for the specific card or component.

Taiwan Battery Recycling Mark



廢電池請回收

Overview

The Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 family of controllers offers redundant array of independent disks (RAID) control capabilities. The PERC 5 Serial Attached Small Computer System Interface (SAS) RAID controllers support SAS devices and Dell-qualified SATA devices. The controllers provide reliability, high performance, and fault-tolerant disk subsystem management.

Scope of the User's Guide

This user's guide for the PERC 5 controllers documents the following topics:

- Basic information about the PERC 5 controllers and RAID functionality
- Information about PERC 5 controller features
- Hardware installation and battery management
- Installation procedures for operating system drivers
- RAID configuration and management
- Troubleshooting information
- Regulatory information and notices

PERC 5 Controller Descriptions

The following list includes a description of each type of controller:

- The PERC 5/E Adapter with two external x4 SAS ports and a transportable battery backup unit (TBBU)
- The PERC 5/i Adapter with two internal x4 SAS ports with or without a battery backup unit, depending on the system
- The PERC 5/i Integrated controller with two internal x4 SAS ports and a battery backup unit

Each controller supports up to 64 virtual disks. In addition, PERC 5/E controllers provide two ports that can connect to up to three enclosures each, for a total of six enclosures per controller. Each enclosure can contain up to 15 physical disks, meaning a controller can support up to 90 physical disks total in the six enclosures.



NOTE: PERC 5/i is limited by the configuration supported on the platform.

PCI Architecture

PERC 5 controllers support a PCI-E x8 host interface. PCI-E is a high-performance I/O bus architecture designed to increase data transfers without slowing down the Central Processing Unit (CPU). PCI-E goes beyond the PCI specification and is intended as a unifying I/O architecture for various systems such as, desktops, workstations, mobiles, server, communications, and embedded devices.

Operating System Support

The PERC 5 controllers support the following operating systems:

- Microsoft® Windows® 2000 Server family
- Windows Server® 2003 (includes Standard, Enterprise, and Small Business Servers)
- Windows Server 2003 DataCenter
- Windows XP
- Windows Vista™
- Red Hat® Enterprise Linux® 3, Red Hat Enterprise Linux 4, and Red Hat Enterprise Linux 5
- SUSE® Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10



NOTE: See driver release on www.dell.com for specific operating system service pack requirements.



NOTE: Microsoft Windows XP is supported with a PERC 5 controller only when the controller is installed in a Dell Precision™ workstation.



NOTE: See the system documentation located at the Dell Support website at support.dell.com for the latest list of supported operating systems and driver installation instructions.

RAID Description

RAID is a group of multiple independent physical disks that provide high performance by increasing the number of drives used for saving and accessing data. A RAID disk subsystem improves input/output (I/O) performance and data availability. The physical disk group appears to the host system as either a single storage unit or multiple logical units. Data throughput improves because several disks are accessed simultaneously. RAID systems also improve data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining data or parity physical disks.



NOTICE: In the event of a physical disk failure, you cannot rebuild data on a RAID 0 virtual disk.

Summary of RAID Levels

RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy.

RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. This is good for small databases or other applications that require small capacity, but complete data redundancy.

RAID 5 uses disk striping and parity data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access.

RAID 10, a combination of RAID 0 and RAID 1, uses disk striping across mirrored spans. It provides high data throughput and complete data redundancy.

RAID 50, a combination of RAID 0 and RAID 5, uses distributed data parity and disk striping and works best with data that requires high system availability, high request rates, high data transfers, and medium-to-large capacity.

RAID Terminology


Disk Striping

Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each physical disk storage space into stripes that can vary in size ranging from 8 KB to 128 KB, often referred to as stripe size. These stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a strip.

For example, in a four-disk system using only disk striping (used in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple physical disks are accessed simultaneously, but disk striping does not provide data redundancy.

Figure 1-1 shows an example of disk striping.

Figure 1-1. Example of Disk Striping (RAID 0)



Disk Mirroring

With mirroring (used in RAID 1), data written to one disk is simultaneously written to another disk. If one disk fails, the contents of the other disk can be used to run the system and rebuild the failed physical disk. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the disk are completely written to a second disk, it does not matter if one of the disks fails. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.

Disk mirroring provides 100 percent redundancy, but is expensive because each physical disk in the system must be duplicated. Figure 1-2 shows an example of disk mirroring.



 **NOTE:** Mirrored physical disks improve read performance by read load balance.

Figure 1-2. Example of Disk Mirroring (RAID 1)




Parity


Parity creates a set of redundant data from two or more parent data sets. You can use the redundant data to rebuild one of the parent data sets. Parity data does not fully duplicate the parent data sets but that data can be used to reconstruct the data if lost. In RAID, this method is applied to entire physical disks or stripes across all the physical disks in a physical disk group.

The parity data is distributed across all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity and the data on the remaining physical disks. RAID level 5 combines distributed parity with disk striping, as shown in Figure 1-3. Parity provides redundancy for one physical disk failure without duplicating the contents of entire physical disks. However, parity generation can slow the write process.

Figure 1-3. Example of Distributed Parity (RAID 5)



Note: Parity is distributed across all drives in the array.

 **NOTE:** Parity is distributed across multiple physical disks in the disk group.

PERC 5 Features

This section describes the features of the The Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 family of controllers, such as the configuration options, disk array performance, hardware specifications, redundant array of independent disks (RAID) management utilities, and operating system software drivers.

PERC 5 Controller Features

This section describes the hardware configuration features for the PERC 5 controllers. Table 2-1 compares the configurations for the controllers.


Table 2-1. PERC 5 Controller Comparisons


Specification	PERC 5/E Adapter	PERC 5/i Adapter	PERC 5/i Integrated
RAID Levels	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10, 50
Enclosures per Port	Up to 3 enclosures	N/A	N/A
Ports	2 x4 external wide port	2 x4 internal wide port	2 x4 internal wide port
Processor	Intel® IOP333 I/O processor with Intel XScale Technology	Intel IOP333 I/O processor with Intel XScale Technology	Intel IOP333 I/O processor with Intel XScale Technology
Battery Backup Unit	Yes, Transportable	Yes ^a	Yes
Cache Memory	256 MB DDR2 cache memory size	256 MB DDR2 cache memory size	256 MB DDR2 cache memory size
Cache Function	Write-back, write-through, adaptive read ahead, non-read ahead, read ahead	Write-back, write-through, adaptive read ahead, non-read ahead, read ahead	Write-back, write-through, adaptive read ahead, non-read ahead, read ahead
Maximum Number of Drives per Array	Up to 32 drives per array	Up to 32 drives per array	Up to 32 drives per array
Maximum Number of Arrays per Disk Group	Up to 8 arrays (in a spanned configuration)	Up to 8 arrays (in a spanned configuration)	Up to 8 arrays (in a spanned configuration)

Table 2-1. PERC 5 Controller Comparisons (continued)

Specification	PERC 5/E Adapter	PERC 5/i Adapter	PERC 5/i Integrated
Maximum Number of Virtual Disks per Disk Group	Up to 16 virtual disks per disk group	Up to 16 virtual disks per disk group	Up to 16 virtual disks per disk group
	RAID 0=16	RAID 0=16	RAID 0=16
	RAID 1=16	RAID 1=16	RAID 1=16
	RAID 5=16	RAID 5=16	RAID 5=16
	RAID 10= 1	RAID 10=1	RAID 10=RAID 50=1
	RAID 50=1	RAID 50=1	
Multiple Virtual Disks per Controller	Up to 64 virtual disks per controller	Up to 64 virtual disks per controller	Up to 64 virtual disks per controller
Support for x8 PCI Express host interface	Yes	Yes	Yes
Online Capacity Expansion	Yes	Yes	Yes
Dedicated and Global Hot Spares	Yes	Yes	Yes
Hot Swap Devices Supported	Yes	Yes	Yes
Non-Disk Devices Supported	No	No	No
Mixed Capacity Physical Disks Supported	Yes	Yes	Yes
Hardware Exclusive-OR (XOR) Assistance	Yes	Yes	Yes


^a The PERC 5/i Adapter supports a battery backup unit (BBU) on selected systems only. See the documentation that shipped with the system for additional information.

 **NOTE:** The maximum length of cable that you can use for Serial Attached SCSI (SAS) is 4 meters (13 feet) from port to port. This applies only to external cables.

 **NOTE:** The maximum array size is limited by the maximum number of drives per array (32) and the maximum number of spans per disk group (8), along with the size of the physical drives. This limits the number of spans in RAID 10 to eight, giving a total of 16 drives per virtual disk.

Compatibility With Virtual Disks Created on Existing PERC 5 Controllers

The PERC 5 controllers recognize and use virtual disks created on existing PERC 5 controllers without risking data loss, corruption, redundancy, or configuration loss. Similarly, the virtual disks created on the controllers can be transferred to other PERC 5 controllers.


 **NOTE:** For more information about compatibility, contact your Dell Technical Support Representative.

SMART Technology

The Self-Monitoring Analysis and Reporting Technology (SMART) feature monitors the internal performance of all motors, heads, and physical disk electronics to detect predictable physical disk failures. This feature helps monitor physical disk performance and reliability, and protects the data on the physical disk. When problems are detected on a physical disk, you can replace or repair the physical disk without losing any data.


SMART-compliant physical disks have attributes for which data (values) can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical and electrical failures display some degradation in performance before failure.

There are numerous factors that relate to predictable physical disk failures, such as a bearing failure, a broken read/write head, and changes in spin-up rate. In addition, there are factors related to read/write surface failure, such as seek error rate and excessive bad sectors.

 **NOTE:** See www.t10.org for detailed information about Small Computer System Interface (SCSI) interface specifications and www.t13.org for Serial Attached ATA (SATA) interface specifications.


Background Initialization

Background Initialization (BGI) is a process to correct parity on the virtual disks. BGI is an automated check for media errors in which parity is created and written in the process. BGI does not run on RAID 0 virtual disks.

 **NOTE:** You cannot permanently disable BGI. If you cancel BGI, it automatically restarts within five minutes. See "Stopping Background Initialization" on page 67 for information on stopping BGI.

The background initialization rate is controlled by the storage management software. You must stop an ongoing background initialization before you change the rate, or the rate change will not take effect. After you stop background initialization and change the rate, the rate change will take effect when the background initialization restarts automatically.

 **NOTE:** Unlike initialization of virtual disks, background initialization does not clear data from the physical disks.

 **NOTE:** Consistency Check and Background Initialization perform the same function. The difference between them is that Background Initialization cannot be started manually, while Consistency Check can.

LED Operation

The LED on the physical disk carrier indicates the state of each physical disk. For internal storage, see your system documentation for more information about the blink patterns.

For blink patterns on the Dell PowerVault™ MD1000, see the *Dell PowerVault MD1000 Hardware Owner's Manual*. For blink patterns on the Dell PowerVault MD3000, see the *Dell PowerVault MD3000 Hardware Owner's Manual*.

Disk Roaming

The PERC 5 controllers support moving physical disks from one cable connection or backplane slot to another on the same controller. The controllers automatically recognize the relocated physical disks and logically place them in the proper virtual disks that are part of the disk group. You can perform disk roaming only when the system is shut down.


Perform the following steps to use disk roaming.

- 1 For proper shutdown, turn off the power to the system, physical disks, enclosures, and system components, and then disconnect the power cords from the system.
- 2 Move the physical disks to different positions on the backplane or enclosure.
- 3 Perform a safety check. Make sure the physical disks are inserted properly and perform correctly.
- 4 Turn on the system.

The controller detects the RAID configuration from the configuration data on the physical disks.

Disk Migration

The PERC 5 controllers support migration of virtual disks from one controller to another without taking the target controller offline. However, the source controller must be offline prior to performing the disk migration. The controller can import a virtual disk that is in optimal or degraded state. A virtual disk in offline state cannot be imported.


 **NOTE:** The PERC 5 controllers are not backward compatible with previous SCSI PERC RAID controllers.

When a controller detects a physical disk with a pre-existing configuration, it flags the physical disk as *foreign* and generates an alert indicating that a foreign disk was detected.

Dedicated hot spares are imported as global hot spares. The firmware generates an alert to indicate a change in hot spare configuration.


Perform the following steps to use disk migration.

- 1 Turn off the system that contains the source controller.
- 2 Move the appropriate physical disks from the source controller to the target controller.
The system with the target controller can be running while inserting the physical disks.
- 3 The storage management application will flag the inserted disks as foreign disks.
- 4 Use the storage management application to import the detected foreign configuration.

 **NOTE:** Ensure the complete set of physical disks that form the virtual disk are migrated.

Alarm Alert in Case of Physical Disk Failures

An audible alarm is available on the PERC 5/E Adapter to alert you of key critical and warning events involving the virtual disk or physical disk problems. You can use the Basic Input/Output System (BIOS) Configuration Utility and management application to enable, disable, or silence the on-board alarm tone.

 **NOTE:** See "Audible Alarm Warnings" on page 79 for information about audible alarm codes.

Battery Management

The Transportable Battery Backup Unit (TBBU) is a cache memory module with an integrated battery pack that enables you to transport the cache module with the battery into a new controller. The TBBU protects the integrity of the cached data on the PERC 5/E Adapter by providing backup power during a power outage.

The Battery Backup Unit (BBU) is a battery pack that protects the integrity of the cached data on the PERC 5/i Adapter and PERC 5/i Integrated controllers by providing backup power during a power outage. Unlike the TBBU, the BBU is not directly attached to the cache memory module and is therefore not transportable with the controller.

The TBBU and BBU offer an inexpensive way to protect the data on the memory module. The lithium battery provides a way to store more power in a smaller form factor than previous batteries.

See "Transferring a TBBU Between Controllers" on page 30 for detailed procedures about handling controller cache in case of a controller failure.

Introduction to Write Cache Policy

The cache controller writes a block of data to cache memory, which is much faster than writing to the physical disk. The cache controller sends an acknowledgement of data transfer completion to the host system.

Write-Back versus Write-Through

In write-through caching, the controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction. The controller then writes the cached data to the storage device when system activity is low or when the write buffer approaches capacity.

In write-back caching, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. The cached data is not written to the storage device.

The risk of using write-back cache is that the cached data can be lost if there is a power failure before it is written to the storage device. This risk is eliminated by using a battery backup unit on selected PERC 5 controllers. Refer to Table 2-1 for information on which controllers support a battery backup unit.

Write-back caching has a performance advantage over write-through caching.



NOTE: The default cache setting is write-back caching.



NOTE: Certain data patterns and configurations perform better in a write-through cache policy.

How Firmware Manages Cache

Firmware manages cache based on the condition of the battery. Learn Cycle is a battery calibration operation performed by the controller periodically (approximately every 3 months) to determine the condition of the battery.

Conditions Under Which Write-Back is Employed

Write-back caching is used under all conditions in which the battery is present and in good condition.

Conditions Under Which Write-Through is Employed

Write-through caching is used under all conditions in which the battery is missing or in a low-charge state. Low-charge state is when the battery is not capable of maintaining data for at least 24 hours in the case of a power loss.

Conditions Under Which Forced Write-Back With No Battery is Employed

Write-Back mode is available when the user selects **Force WB with no battery**. When Forced Write-Back mode is selected, the virtual disk is in Write-Back mode even if the battery is present and in good condition (healthy), or if a learn cycle is in process. Dell recommends that you use a power backup system when forcing Write-Back to ensure that there is no loss of data if the system suddenly loses power.

Learn Cycle Completion Time Frame

The time frame for completion of a learn cycle is a function of the battery charge capacity and the discharge/charge currents used. For PERC 5, the expected time frame for completion of a learn cycle is approximately seven hours and consists of the following parts:

- Learn cycle discharge cycle: approximately three hours
- Learn cycle charge cycle: approximately four hours

Learn cycles shorten as the battery capacity derates over time.



NOTE: See the storage management application for additional information.

During the discharge phase of a learn cycle, the PERC 5 battery charger is disabled. In this phase, the battery voltage is monitored through the smart battery bus (SMBus) using the battery gas-gauge. When the battery reaches the discharge capacity threshold (DCT), the dummy load is disabled and the battery charger is re-enabled. At this point, the battery charger detects that the battery is below the fast charge trigger voltage (FCTV) and initiates a fast-charge of the battery. The learn cycle completes once the battery fast charge is completed.

RAID Configuration Information

Table 2-2 lists the configuration features for the PERC 5 controllers.

Table 2-2. Features for RAID Configuration

Specification	PERC 5/E Adapter	PERC 5/i Adapter and 5/i Integrated
Number of virtual disks supported	Up to 64 virtual disks per controller	Up to 64 virtual disks per controller NOTE: The number of physical disks on a controller is limited by the backplane on which the card is attached.
Online RAID level migration	Yes	Yes
Disk roaming	Yes	Yes
No reboot necessary after capacity expansion	Yes	Yes
User-specified rebuild rate	Yes	Yes

Fault Tolerance Features

Table 2-3 lists the features that provide fault tolerance to prevent data loss in case of a failed physical disk.


Table 2-3. Fault Tolerance Features

Specification	Feature
Support for SMART	Yes
Support for Patrol Read	Yes
Physical disk failure detection	Automatic
Physical disk rebuild using hot spares	Automatic
Parity generation and checking (RAID 5 only)	Yes
Battery backup for controller cache to protect configuration data	Yes ^a
Hot-swap manual replacement of a physical disk unit without bringing the system down	Yes


^a The PERC 5/i Adapter supports a battery backup unit (BBU) on selected systems only. For additional information, see documentation that was shipped with the system.

Physical Disk Hot Swapping

Hot swapping is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions).

 **NOTE:** The backplane or enclosure must support hot swapping in order for the PERC 5 controllers to support hot swapping.

 **NOTE:** Ensure that SAS drives are replaced with SAS drives and SATA drives are replaced with SATA drives.

 **NOTE:** While swapping a disk, ensure that the new disk is of equal or greater capacity than the disk that is being replaced.

Failed Physical Disk Detection

The firmware automatically detects and rebuilds failed physical disks. Automatic rebuilds can be performed transparently with hot spares. If you have configured hot spares, the controllers automatically try to use them to rebuild failed physical disks.

Patrol Read

The Patrol Read function is designed as a preventive measure that includes review of your system for possible physical disk errors that could lead to physical disk failure and damage data integrity. The Patrol Read operation can find and possibly resolve any potential problem with physical disks prior to host access. This can enhance overall system performance because error recovery during a normal I/O operation may not be necessary. You can use the storage management application to perform Patrol Read functions.

Patrol Read Behavior

The following is an overview of Patrol Read behavior:

- 1 Patrol Read runs on all disks on the controller that are configured as part of a virtual disk including hot spares. Patrol Read does not run on unconfigured physical disks. Unconfigured disks are those that are not part of a virtual disk or are in a ready state.
- 2 Patrol Read adjusts the amount of controller resources dedicated to Patrol Read operations based on outstanding disk I/O. For example, if the system is busy processing I/O operation, then Patrol Read will use fewer resources to allow the I/O to take a higher priority.
- 3 Patrol Read operates on all configured physical disks on the controller and there is no method to deselect disks.
- 4 In a Patrol Read iteration, Patrol Read will restart from zero percent if in **Auto** mode. In **Manual** mode, Patrol Read does not restart on a reboot. Use **Manual** mode if you have selected a window of time dedicated to running Patrol Read.

Configuration

You can use the storage management application to select the Patrol Read options. Use Patrol Read options to set automatic or manual operation, or disable Patrol Read. The following sections describe Patrol Read functions and operations available in the Storage Management application.



NOTE: See the storage management application's documentation for more information about the Patrol Read configuration features available.

Patrol Read Modes

The following describes the scheduling details for Patrol Read:

- 1 By default the controller sets the Patrol Read to **Auto** mode. You can set the Patrol Read to either **Auto** or **Manual** mode.
- 2 In **Auto** mode, Patrol Read runs continuously on the system and is scheduled to start a new Patrol Read within seven days after the last iteration is completed.
- 3 When **Patrol Read Mode** is changed from **Auto** to **Manual**, or **Auto** to **Disabled**, the **Next execution will start at:** field will be set to **N/A**.

Behavior Details

The behavior details of Patrol Read are as follows:

- 1 Setting Patrol Read in **Manual** mode does not start Patrol Read. It only sets the mode so that you can select **Start** whenever you want to run Patrol Read. When the mode is set to **Manual**, it remains in that mode until you change it.
- 2 Setting the mode to **Automatic** starts Patrol Read. When the Patrol Read operation is complete, it will set itself to run within seven days of the last iteration.

Blocked Operations


If any of the following conditions exist, then Patrol Read will not run on any of the affected disks:

- An unconfigured disk (the disk is in the **READY** state)
- Disks that are members of a virtual disk undergoing a reconstruction
- Disks that are members of a virtual disk undergoing a Background Initialization or Consistency Check


Hardware Installation and Configuration

This chapter describes how to install the Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 cards.

Installing the PERC 5/E Adapter


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

- 1 Unpack the PERC 5/E Adapter and check for damage.

 **NOTE:** Contact Dell Technical support if the controller is damaged.

- 2 Turn off the system and attached peripherals, and disconnect the system from the electrical outlet. See your system's *Hardware Owner's Manual* for more information on power supplies.
- 3 Disconnect the system from the network and remove the cover of the system. See your system's *Hardware Owner's Manual* for more information on opening the system.
- 4 Select an empty PCI Express (PCI-E) slot. Remove the blank filler bracket on the back of the system aligned with the PCI-E slot you have selected.
- 5 Align the PERC 5/E Adapter to the PCI-E slot you have selected.
- 6 Insert the controller gently, but firmly, until the controller is firmly seated in the PCI-E slot. See Figure 3-1.

 **CAUTION:** Never apply pressure to the memory module while inserting the controller into the PCI-E slot. Applying pressure could break the module.

 **CAUTION:** Electrostatic discharge can damage sensitive components. Always use proper antistatic protection when handling components. Touching components without using a proper ground can damage the equipment.




 **NOTE:** See your system documentation for a list of compatible controllers.

Figure 3-1. Installing a PERC 5/E Adapter



- 7 Tighten the bracket screw, if any, or use the system's retention clips to secure the controller to the system's chassis.
- 8 Replace the cover of the system. See your system's *Hardware Owner's Manual* for more information on closing the system.
- 9 Connect the cable from the external enclosure to the controller. See Figure 3-2.


Figure 3-2. Connecting the Cable From the External Enclosure





- 10 Reconnect the power cable(s) and network cables, and then turn on the system.

Installing the Transportable Battery Backup Unit (TBBU) for PERC 5/E


This section describes the installation of the transportable battery backup unit (TBBU) on the PERC 5/E.


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


 **NOTE:** All work must be performed at an Electrostatic Discharge (ESD)-safe workstation to meet the requirements of EIA-625 - "Requirements For Handling Electrostatic Discharge Sensitive Devices." All actions must be performed following the IPC-A-610 latest revision ESD recommended practices.

 **NOTE:** Charge the PERC 5 battery before initial use to attain full functionality.

- 1 Unpack the TBBU and follow all antistatic procedures.

 **CAUTION:** When unpacking static sensitive component from its shipping carton, do not remove the component from the antistatic packing material until you are ready to install the component. Just before unwrapping the antistatic package, ensure that you discharge static electricity from your body.

 **NOTICE:** When transporting a sensitive component, first place it in an antistatic container or packaging.

 **NOTE:** Handle all sensitive components in a static-safe area. If possible, use antistatic floor pads and work bench pads.

- 2 With the DIMM removed from the controller, insert one end of the battery pack harness (the red, white, yellow, and green wires) into the connector on the memory module and the other end into the connector on the battery.
- 3 Place the top edge of the battery over the top edge of the memory module so that the arms on the side of the battery fit into their sockets on the memory module. See Figure 3-3.



 **CAUTION:** Electrostatic discharge can damage sensitive components. Always use proper antistatic protection when handling components. Touching components without using a proper ground can damage the equipment.

Figure 3-3. Installing a TBBU



- 4 Place the PERC 5/E on a flat, clean, static-free surface.
- 5 Mount the memory module in the controller memory socket like a standard DIMM. See "Installing the DIMM on a PERC 5/E Adapter" on page 28 for more information.

The memory module is mounted flush with the board so that the memory module is parallel to the board when installed.

- 6 Press the memory module firmly into the memory socket.

As you press the memory module into the socket, the TBBU clicks into place, indicating that the controller is firmly seated in the socket, and the arms on the socket fit into the notches to hold the memory module securely.

Installing the DIMM on a PERC 5/E Adapter

This section describes how to install the memory module on a PERC 5/E Adapter.

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

🔄 NOTICE: PERC 5 cards support DELL-qualified 256-MB DDRII 400MHz ECC-registered DIMMs with x16 DRAM components. Installing unsupported memory causes the system to hang at POST.

- 1 Remove the memory module in an antistatic environment.

🔧 NOTE: When unpacking a static sensitive component from its shipping carton, do not remove the component from the antistatic packing material until you are ready to install the component. Just before unwrapping the antistatic package, ensure to discharge static electricity from your body.

🔧 NOTE: Handle all sensitive components in a static-safe area. If possible, use antistatic floor pads and work bench pads.

🔧 NOTE: Do not touch the gold leads and do not bend the memory module.

- 2 Align the memory module so that the keyed edge of the memory module is placed exactly on top of the physical divider on the memory socket of the controller to avoid damage to the DIMM.
- 3 Insert the memory module on the memory socket of the controller and apply a smooth, downward pressure on both ends or on the middle of the memory module until the retention clips fall into the allotted slots on either side of the memory module. See Figure 3-4.




 **NOTE:** Figure 3-4 displays the installation of a memory module on a PERC 5/E Adapter.

Figure 3-4. Installing a DIMM



Transferring a TBBU Between Controllers

The TBBU provides uninterrupted power supply to the memory module for up to 72 hours if power supply is unexpectedly interrupted while cached data is still present. If the controller fails as a result of a power failure, you can move the TBBU to a new controller and recover the data. The controller that replaces the failed controller must be devoid of any prior configuration.


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

Perform the following steps to replace a failed controller with data in the TBBU:

- 1 Turn off the system and disconnect the physical disks.
- 2 Remove the controller that has the TBBU currently installed from the system.
- 3 Remove the TBBU from the controller.
- 4 Insert the TBBU into the new controller.
See "Installing the Transportable Battery Backup Unit (TBBU) for PERC 5/E" on page 27.
- 5 Insert the new controller into the system.
See the relevant sections on installing controllers under "Installing the PERC 5/E Adapter" on page 25.
- 6 Turn on the system.
The controller flushes the cache data to the virtual disks.


Removing the PERC 5/E Adapter

This section describes how to remove the PERC 5/E from a system.

 **NOTE:** In the event that the SAS cable is accidentally pulled when the system is operational, reconnect the cable and refer to the online help of your storage management application for the required recovery steps.


- 1 Perform a controlled shutdown on the system in which the PERC 5/E is installed, as well as any attached storage enclosures.
- 2 Disconnect the system from the electrical outlet and remove the system cover.

 **CAUTION:** Running a system without the system cover installed may cause damage due to improper cooling.

 **NOTE:** For more information on removing and reinstalling the system cover, see the *Hardware Owner's Manual* that shipped with the system.

- 3 Locate the PERC 5/E in the system and disconnect the external cables from the PERC 5/E.
- 4 Remove any retention mechanism such as a bracket screw that may be holding the PERC 5/E in the system and gently lift the controller from the system's PCI-E slot. See Figure 3-5.

Figure 3-5. Removing the PERC 5/E Adapter



NOTE: For more information on removing peripherals installed in the system's PCI-E slots, see the *Hardware Owner's Manual* that shipped with the system.


Removing the DIMM and Battery from a PERC 5/E Adapter

NOTE: The TBBU on the PERC 5/E Adapter consists of a DIMM and battery backup unit.

This section describes how to remove the TBBU from a PERC 5/E Adapter that is currently installed in a system.

- 1 Perform a controlled shutdown on the system in which the PERC 5/E Adapter is installed, as well as any attached storage enclosures, and remove the PERC 5/E Adapter from the system by following the instructions detailed in "Removing the PERC 5/E Adapter" on page 30.
- 2 Visually inspect the controller and determine whether the dirty cache LED on the DIMM is illuminated. See Figure 3-6. If the LED is illuminated, reinsert the controller into the system, replace the system cover, reconnect the system to power, turn the system on and repeat step 1.


Figure 3-6. PERC 5/E Adapter Dirty Cache LED Location



⚠ CAUTION: Running a system without the system cover installed may cause damage due to improper cooling.

- 3** Remove the TBBU assembly from the adapter by pressing down on the tabs at each edge of the DIMM connector and lifting the TBBU off the adapter.
- 4** Disconnect the battery cable from the DIMM.
- 5** Detach the battery from the DIMM by pressing out on the battery clips inserted through the DIMM and rotating the battery off the DIMM. See Figure 3-7.

Figure 3-7. Removing the TBBU



Installing the PERC 5/i Adapter

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

🔧 NOTE: See your system documentation for information on installing the PERC 5/i.

🔧 NOTE: The PERC 5/i Adapter will have a battery backup unit (BBU) when installed on a PowerEdge server and will not have one when installed on a Dell Precision™ workstation or a PowerEdge SC server.

- 1 Unpack the PERC 5/i Adapter and check for damage.

🔧 NOTE: Call Dell Technical support if the controller is damaged.

- 2 Turn off the system and attached peripherals, and disconnect the system from the electrical outlet. See your system's *Hardware Owner's Manual* for more information on power supplies.

- 3 Disconnect the system from the network and remove the cover of the system.

🔧 NOTE: For more information on removing and replacing the system cover, see the *Hardware Owner's Manual* included with the system.


- 4 Select an empty PCI-E slot and remove the blank filler bracket. The blank filler bracket is located on the back of the system, aligned with the empty PCI-E slot you have selected.

- 5 Align the PERC 5/i Adapter to the PCI-E slot you have selected.
- 6 Insert the adapter gently but firmly until the board is firmly seated in the PCI-E slot. See Figure 3-8.

⚠ CAUTION: Do not apply pressure to the memory module while inserting the controller into the PCI-E slot. Applying pressure could break the memory module or damage the DIMM connector.


⚠ CAUTION: Electrostatic discharge can damage sensitive components. Always use proper antistatic protection when handling components. Touching components without using a proper ground can damage the equipment.

Figure 3-8. Installing the PERC 5/i Adapter




- 7 Tighten the bracket screw, if any, or use the system's retention clips to secure the controller to the system's chassis.
- 8 Connect the cables from the backplane of the system to the controller. See Figure 3-9.

Figure 3-9. Connecting Cables to the Controller




- 9 Replace the cover of the system. See your system's *Hardware Owner's Manual* for more information on closing the system.
- 10 Reconnect the power cable(s) and network cables, and then turn on the system.


Removing the PERC 5/i Adapter

 **NOTE:** A PERC 5/i Adapter installed on a Dell Precision workstation or PowerEdge SC server does not have a BBU.

This section describes how to remove the PERC 5/i Adapter or PERC 5/i while it is installed in a system.

- 1 Perform a controlled shutdown on the system in which the PERC 5/i is installed.
- 2 Disconnect the system from the electrical outlet, and remove the cover of the system.

 **CAUTION:** Running a system without the system cover installed may cause damage due to improper cooling.


 **NOTE:** For more information on removing and reinstalling the system cover, see the *Hardware Owner's Manual* that shipped with the system.

- 3 Determine whether the dirty cache LED on the controller is illuminated.
 - If the LED is illuminated, replace the system cover, reconnect the system to power, turn on the system, and repeat step 1 and step 2. See Figure 3-10.

NOTE: The location of the PERC 5/i varies from system to system. See the *Hardware Owner's Manual* included with the system for specific information on where the PERC 5/i is located.


- If the LED is not illuminated, continue with the next step.

Figure 3-10. PERC 5/i Adapter Dirty Cache LED Location



- 4 Disconnect the data cables and battery cable from the PERC 5/i.
- 5 Remove any retention mechanism, such as a bracket screw, that might be holding the PERC 5/i in the system, and gently lift the controller from the system's PCI-E slot. See Figure 3-11.

Figure 3-11. Removing the PERC 5/i Adapter



NOTE: See the *Hardware Owner's Manual* for information on removing the PERC 5/i from the system.

Removing the DIMM from a PERC 5/i Card

This section describes how to remove the memory module on a PERC 5/i Adapter or PERC 5/i that is currently installed in a system.

NOTE: PERC 5 cards support DELL-qualified 256-MB DDRII 400MHz ECC-registered DIMMs with x16 DRAM components. Installing unsupported memory causes the system to hang at POST.


- 1 Perform a controlled shutdown on the system in which the PERC 5/i is installed and remove the PERC 5/i from the system by following the instructions detailed in the section "Removing the PERC 5/i Adapter" on page 35.

CAUTION: Running a system without the system cover installed can cause damage due to improper cooling.

- 2 Remove the DIMM by pressing down on the tabs at each edge of the DIMM connector and lifting the DIMM off the controller.

NOTE: The location of the PERC 5/i varies from system to system. See the *Hardware Owner's Manual* that shipped with the system for specific information on the location of the PERC 5/i.


Disconnecting the BBU from a PERC 5/i Adapter or a PERC 5/i

 **NOTE:** A PERC 5/i Adapter installed in a Dell Precision workstation or a PowerEdge SC server does not have a BBU.


This section describes how to disconnect the BBU on a PERC 5/i Adapter or a PERC 5/i while the adapter or controller is installed in a system.

- 1 Perform a controlled shutdown on the system in which the PERC 5/i is installed.
- 2 Disconnect the system from the electrical outlet and remove the system cover.

 **CAUTION:** Running a system without the system cover installed can cause damage due to improper cooling.

 **NOTE:** For more information on removing and replacing the system cover, see the *Hardware Owner's Manual* that shipped with the system.

- 3 Determine whether the dirty cache LED on the controller is illuminated.
 - If the LED is illuminated, replace the system cover, reconnect the system to power, turn on the system, and repeat step 1 and step 2. See Figure 3-10.

 **NOTE:** The location of the PERC 5/i varies from system to system. See the *Hardware Owner's Manual* that shipped with the system for specific information on where the PERC 5/i is located.
 - If the LED is not illuminated, continue with the next step.
- 4 Locate the battery cable connection next to the DIMM on the controller, and disconnect the battery.

Driver Installation

The Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 family requires software drivers to operate with the Microsoft® Windows®, Red Hat® Enterprise Linux®, and SUSE® Linux operating systems.

This chapter contains the procedures for installing the drivers for the following operating systems.

- Microsoft® Windows® 2000 Server family
- Windows Server® 2003 (includes Standard, Enterprise, and Small Business Servers)
- Windows Server 2003 DataCenter
- Windows XP
- Windows Vista™
- Red Hat Enterprise Linux 3, 4, and 5
- SUSE Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10



NOTE: See the Dell Support website at support.dell.com to check operating system compatibility.

The two methods for installing a driver that are discussed in this chapter are:

- During operating system installation. Use this method if you are performing a new installation of the operating system and want to include the drivers.
- Updating existing drivers. Use this method if the operating system and PERC 5 controller are already installed and you want to update to the latest drivers.

Creating a Driver Diskette

Perform the following steps to create a driver diskette.

- 1 Browse to the download section for the system from the Dell Support website at support.dell.com.
- 2 Locate and download the latest PERC 5 controller driver to the system. The driver should be labeled as packaged for a diskette on the Dell Support website.
- 3 Follow the instructions on the Dell Support website for extracting the driver to the diskette.


Installing the Driver During a Microsoft Operating System Installation

Perform the following steps to install the driver during operating system installation.

- 1 Boot the system using the Microsoft Windows 2000/Windows XP/Microsoft Windows Server 2003 CD.
- 2 When the message **Press F6 if you need to install a third party SCSI or RAID driver** appears, press the <F6> key immediately.

Within a few minutes, a screen appears that asks for additional controllers in the system.

- 3 Press the <S> key.
The system prompts for the driver diskette to be inserted.
- 4 Insert the driver diskette in the diskette drive and press <Enter>.
A list of PERC controllers appears.
- 5 Select the right driver for the installed controller and press <Enter> to load the driver.

 **NOTE:** For Windows Server 2003, a message may appear that states that the driver that you provided is older or newer than the existing Windows driver. Press <S> to use the driver that is on the floppy diskette.

- 6 Press <Enter> again to continue the installation process as usual.

Installing a Windows Server 2003, Windows Vista, Windows 2000, or XP Driver for a New RAID Controller

Perform the following steps to configure the driver for the RAID controller on a system that already has Windows installed.


 **NOTE:** Windows Vista™ has PERC 5 driver native on the Windows Vista installation CD. For the latest supported drivers, go to the Dell Support website at support.dell.com.

- 1 Turn off the system.
- 2 Install the new RAID controller in the system.
See "Hardware Installation and Configuration" on page 25 for detailed instructions on installing and cabling the RAID controller in the system.
- 3 Turn on the system.
The Windows operating system detects the new controller and displays a message to inform the user.
- 4 The **Found New Hardware Wizard** screen pops up and displays the detected hardware device.
- 5 Click **Next**.
- 6 On the **Locate device driver** screen, select **Search for a suitable driver for my device** and click **Next**.
- 7 Insert the appropriate driver diskette and select **Floppy disk drives** on the **Locate Driver Files** screen.
- 8 Click **Next**.

- 9 The wizard detects and installs the appropriate device drivers for the new RAID controller.
- 10 Click **Finish** to complete the installation.
- 11 Reboot the server.

Updating an Existing Windows 2000, Windows Server 2003, Windows XP, or Windows Vista Driver

Perform the following steps to update the Microsoft Windows driver for the PERC 5 controller already installed on your system.

 **NOTE:** It is important that you close all applications on your system before you update the driver.

- 1 Select **Start**→ **Settings**→ **Control Panel**→ **System**.

The **System Properties** screen appears.

 **NOTE:** For systems running a Microsoft Windows Server 2003 operating system, click **Start**→ **Control Panel**→ **System**.

- 2 Click on the **Hardware** tab.

- 3 Click **Device Manager**.

The **Device Manager** screen appears.

- 4 Double-click on **SCSI and RAID Controllers**.

 **NOTE:** In Windows Vista, PERC is listed under **Storage Controllers**.

- 5 Double-click the RAID controller for which you want to update the driver.

- 6 Click the **Driver** tab and click **Update Driver**.

The **Upgrade Device Driver Wizard** screen appears.


- 7 Insert the appropriate driver diskette.

- 8 Select **Install from a list or specific location**.

- 9 Click **Next**.

- 10 Follow the steps in the Wizard to search the diskette for the driver.

- 11 Select the INF file from the diskette.

 **NOTE:** For systems running a Windows Server 2003 operating system, select the name of the driver, not the INF file.

- 12 Click **Next** and continue the installation steps in the Wizard.

- 13 Click **Finish** to exit the wizard and reboot the system for the changes to take place.

Installing Windows Vista for a New RAID Controller

Perform the following steps to configure the driver when you add the RAID controller to a system that already has Windows Vista installed.

- 1 Insert the Windows Vista DVD and reboot the system.
- 2 Select location information, such as your country and language.
- 3 Enter the Windows key.
- 4 Select the drive you want to install Windows on by highlighting your choice.
- 5 Select **Load Driver** if needed.



NOTE: You do not need this for PERC or SAS installation as the driver is native on Windows Vista.

The driver can be provided by CD, diskette, DVD, or USB flash driver.

- 6 Browse to the proper location and select the driver.
- 7 Select **Next** if selected unallocated space is the target location.
- 8 Select **Advance** to create a partition for the operating system.
If a partition is manually created, it should at least be 10 GB.



NOTE: See the Windows Vista user documentation to determine the partition size.

- 9 Select **Next**.

The operating system installation starts on the highlighted choice (partition or unallocated volume) and the system will reboot several times.

Do *not* press any keys to start from the CD/DVD.

- 10 Enter user and password information.
- 11 Enter the computer name.
- 12 Select the automatic Windows protection level.
- 13 Set the time, date, and time zone.
- 14 Click **Start**.

Windows checks your system and prompts you for the login password after a few minutes.

Installing the Red Hat Enterprise Linux Driver

Use the procedures in this section to install the driver for Red Hat Enterprise Linux AS (versions 3 and 4) and ES (versions 3 and 4). The driver is updated frequently. To ensure that you have the current version of the driver, download the updated Red Hat Enterprise Linux driver from the Dell Support website at support.dell.com.

Creating a Driver Diskette

Before beginning the installation, copy the drivers from the *Service and Diagnostic Utilities* CD or download the driver appropriate for Red Hat Enterprise Linux (versions 3 and 4) from the Dell Support website at support.dell.com to your temporary directory. This file includes two Red Hat Package Managers (RPMs) and driver update disk files. The package also contains the Dynamic Kernel Module Support (DKMS) Red Hat Package Manager (RPM) file.

The package is a gzipped tar file. After downloading the package to a Linux system, perform the following steps.

- 1 Execute `gunzip` on the package.
- 2 Execute `tar xvf` on the package.

Note that the package contains DKMS RPM, the driver RPM (dkms enabled) and the Driver Update Diskette (DUD) image(s).

- 3 Use the `dd` command to create a driver update disk. Use the appropriate image for the purpose.
`dd if=<name of the dd image file> of=/dev/fd0`
- 4 Use the diskette for operating system installation as described later in this section.

File and Directories Needed to Create the Driver Update Diskette (DUD)

The following files are needed before you create the DUD.



NOTE: The `megaraid_sas` driver package installs these files. You do not need to do anything at this point.

- 1 There is a directory `/usr/src/megaraid_sas-<driver_version>`, which contains the driver source code, `dkms.conf`, and specifications file for the driver.
- 2 In this directory, there is a subdirectory called `redhat_driver_disk` which contains the files needed to create the DUD. The files needed are `disk_info`, `modinfo`, `modules.dep`, and `pcitable`.
- 3 To create the DUD image for pre-Red Hat Enterprise Linux 4, the kernel source package must be installed to compile the driver. For Red Hat Enterprise Linux 4 distribution, the kernel source is not needed.

Creating a Driver Update Diskette

Perform the following steps to create the DUD using the DKMS tool:

- 1 Install the DKMS-enabled `megaraid_sas` driver rpm package on a system running a Red Hat Enterprise Linux operating system.

- 2 Type the following command in any directory:


```
dkms mkdriverdisk -d redhat -m megaraid_sas -v <driver version>
-k <kernel version>
```

This starts the process to create the `megaraid_sas` DUD image.

- 3 If you want to build the DUD image for multiple kernel versions, use the following command:

```
dkms mkdriverdisk -d redhat -m megaraid_sas -v <driver version>
-k <kernel_version_1>, <kernel_version_2>, ...
```

- 4 After the DUD image has been built, you can find it in the DKMS tree for the `megaraid_sas` driver.

 **NOTE:** The DKMS package supports both Red Hat Enterprise Linux and SUSE Linux.

Installing the Driver

Perform the following steps to install Red Hat Enterprise Linux (versions 3 and 4) and the appropriate driver.

- 1 Boot normally from the Red Hat Enterprise Linux installation CD.

- 2 At the command prompt, type:

```
Linux expert dd
```

- 3 When the install prompts for a driver diskette, insert the diskette and press <Enter>.

See "Creating a Driver Diskette" on page 39 for information about creating a driver diskette.

- 4 Complete the installation as directed by the installation program.

The driver will be installed.

Installing the Driver Using an Update RPM

The following steps explain the installation procedure for the Red Hat Enterprise Linux, AS (versions 3 and 4) operating systems and the appropriate RAID driver using an update RPM.

Installing the RPM Package With DKMS Support

Perform the following steps to install the RPM package with DKMS support:

- 1 Uncompress the zipped file of the DKMS-enabled driver package.
- 2 Install the DKMS package using the command `rpm -Uvh <DKMS package name>`.
- 3 Install the driver package using the command `rpm -Uvh <Driver package name>`.
- 4 Reboot the system to load the new driver.

Upgrading the Kernel

If you run the `up2date` utility to upgrade the kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update the kernel:

- 1 In a terminal window, type the following:

```
dkms build -m <module_name> -v <module version> -k <kernel version>
```

```
dkms install -m <module_name> -v <module version> -k <kernel version>
```

- 2 To check whether the driver is successfully installed in the new kernel, type:

```
dkms status
```

You must see a message similar to the following one on the screen to confirm installation:

```
<driver name>, <driver version>, <new kernel version>: installed
```

Installing the SUSE Linux Enterprise Server (Version 9 or 10) Driver

Use the procedures in this section to install the driver for SUSE Linux Enterprise Server (version 9 or version 10). The driver is updated frequently. To ensure you have the current version of the driver, download the updated SUSE Linux Enterprise Server (version 9 or 10) driver from the Dell Support website at support.dell.com.

Installing the Driver Using an Update RPM

The following procedure explains the installation procedure for the SUSE Linux Enterprise Server (version 9 or 10) operating system and the appropriate RAID driver using an update RPM.

Installing the RPM Package With DKMS Support

Perform the following steps to install the RPM package with DKMS support.

- 1 Uncompress the zipped file of the DKMS-enabled driver package.
- 2 Install the DKMS package using the command `rpm -Uvh <DKMS package name>`.
- 3 Install the driver package using the `rpm -Uvh <Driver package name>`.
- 4 Reboot the system to load the new driver.

File and Directories Needed to Create the Driver Update Diskette (DUD)

The following files are needed before you create the DUD.



NOTE: The `megaraid_sas` driver package installs these files. You do not need to do anything at this point.

- 1 The directory `/usr/src/megaraid_sas-<driver_version>` contains the driver source code, `dkms.conf`, and specifications file for the driver.
- 2 In this directory, the subdirectory `redhat_driver_disk` contains the files needed to create the DUD. The files needed are `disk_info`, `modinfo`, `modules.dep`, and `pcitable`.

DUD Creation Procedure

Perform the following steps to create the DUD using the DKMS tool.

- 1 Install the DKMS-enabled megaraid_sas driver rpm package on a system running a SUSE Linux operating system.

- 2 Type the following command in any directory:


```
dkms mkdriverdisk -d redhat -m megaraid_sas -v <driver version>
-k <kernel version>
```

This starts the process to create the megaraid_sas DUD image.


- 3 If you want to build the DUD image for multiple kernel versions, use the following command:

```
dkms mkdriverdisk -d redhat -m megaraid_sas -v <driver version>
-k <kernel_version_1>, <kernel_version_2>, ...
```

- 4 After the DUD image has been built, you can find it in the DKMS tree for the megaraid_sas driver.

 **NOTE:** The DKMS package supports both Red Hat Enterprise Linux and SUSE Linux Enterprise Server.

Installing SUSE Linux Enterprise Server Using the Driver Update Diskette

 **NOTE:** See "Creating a Driver Diskette" on page 39 for information about creating a driver diskette.

Perform the following steps to install SUSE Linux Enterprise Server (version 9 or 10) using the DUD. See "Creating a Driver Diskette" on page 39 for information about creating a driver diskette.

- 1 Insert the appropriate SUSE Linux Enterprise Server (version 9 or 10) Service Pack (SP) CD in the system.
- 2 Select **Installation** from the menu options.
- 3 Select the following key for the driver update disk, depending on the SUSE Linux Enterprise Server version you are installing:
 - a For SUSE Linux Enterprise Server 9, select <F6>.
 - b For SUSE Linux Enterprise Server 10, select <F5>.

- 4 Press <Enter> to load the Linux kernel.

- 5 At the prompt `Please insert the driver update floppy/CDROM`, click **OK**.

The system selects the driver from the diskette and installs it. The system then displays the message `DRIVER UPDATE ADDED` with the description of the driver module.

- 6 Click **OK**.

If you want to install from another driver update medium, continue with the following steps.

- 7 The system displays the message `PLEASE CHOOSE DRIVER UPDATE MEDIUM`.

- 8 Select the appropriate driver update medium.

The system selects the driver from the diskette and installs it.

Upgrading the Kernel

If you run the `up2date` utility to upgrade the kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update the kernel:

- 1 Type the following in a terminal window:

```
dkms build -m <module_name> -v <module version> -k <kernel version>
```

```
dkms install -m <module_name> -v <module version> -k <kernel version>
```

- 2 To check whether the driver is successfully installed in the new kernel, type:

```
dkms status
```

You must see a message similar to the following one on the screen to confirm installation:

```
<driver name>, <driver version>, <new kernel version>: installed
```


RAID Configuration and Management

Storage management applications enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance. The applications for Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 controllers include:

- Dell™ OpenManage™ Storage Management
- SAS RAID Storage Manager
- BIOS Configuration Utility (Ctrl-R)

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for PowerEdge™ servers that provides enhanced features for configuring a system's locally-attached RAID and non-RAID disk storage. Dell OpenManage Storage Management enables you to perform controller and enclosure functions for all supported RAID and non-RAID controllers and enclosures from a single graphical or command-line interface without requiring the use of the controller BIOS utilities. The graphical user interface (GUI) is wizard-driven with features for novice and advanced users and detailed online help. The command line interface is fully-featured and scriptable. Using Dell OpenManage Storage Management, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks.

SAS RAID Storage Manager

SAS RAID Storage Manager is a storage management application for PowerEdge™ SC servers and Dell Precision™ workstations. SAS RAID Storage Manager configures virtual disks, and monitors and maintains PERC 5 controllers, battery backup units, and other devices running on systems and workstations. It offers a graphical user interface (GUI) you can use to perform these tasks.


BIOS Configuration Utility

The BIOS Configuration Utility, also known as Ctrl-R, is a storage management application embedded on the PERC 5 controllers that configures and maintains RAID disk groups and virtual disks, and manages the RAID system. It is independent of any operating system.



NOTE: The BIOS Configuration Utility is to be used for initial setup and disaster recovery. Advanced features can be set through Storage Management and SAS RAID Storage Manager.

The following sections will help you attain the basic knowledge needed to use the BIOS Configuration Utility. See the online help option by pressing <F1> to obtain additional information about the ongoing operation.

 **NOTE:** PERC 5 controller configuration utility refreshes the screen to show changes to information on the screen. The refresh occurs when you press a key or every 15 seconds if you do not press a key.

Entering the BIOS Configuration Utility

The BIOS Configuration Utility configures physical disk groups and virtual disks. Because the utility resides in the controller BIOS, its operation is independent of the operating systems on your system.

Starting the BIOS Configuration Utility


Perform the following steps to enter the BIOS Configuration Utility when you boot the system.

- 1 Turn on and boot the system.

A BIOS banner displays information about the controller and configuration.

- 2 During bootup, press <Ctrl><R> when prompted by the BIOS banner.

After you press <Ctrl><R>, the **Virtual Disk Management** menu screen displays by default. This screen displays the controllers, virtual disks, physical disks, free space, hot spares, and other details in an expandable tree directory similar to that of Microsoft® Windows® Explorer.

 **NOTE:** You can access multiple controllers through the BIOS Configuration Utility by pressing <F12>. Verify which controller you are currently set to edit.

Exiting the Configuration Utility

To exit the BIOS Configuration Utility, press <Esc> at any menu screen until the utility closes.

Menu Navigation Controls


Table 5-1 displays the menu keys you can use to move between the different screens in the BIOS Configuration Utility.

Table 5-1. Menu Navigation Keys


Notation	Meaning and Use	Example
→	Use the right arrow key to open a submenu, move from a menu heading to the first submenu, or move to the first item in that submenu. If you press the right arrow key at a menu heading, the submenu will expand. Press it again to go to the first item in the submenu. The right arrow key is also used to close a menu list in a popup window. Word wrap is supported.	Start → Programs
←	Use the left arrow key to close a submenu, move from a menu item to the menu heading for that item, or move from a submenu to a higher level menu. If you press the left arrow key at a menu heading, the submenu will collapse. Press it again to go to the higher-level menu. Word wrap is supported.	Controller 0 ← Disk Group 1
↑	Use the up arrow key to move to the upper menu items within a menu or to a higher level menu. You can also use the up arrow key to close a menu list in a popup window, such as the stripe element size menu. Word wrap is supported.	Virtual Disk 1 ↑ Virtual Disk 4
↓	Use the down arrow key to move to the lower menu items within a menu or to a lower level menu. You can also use the down arrow key to open a menu list in a popup window, such as the stripe element size menu, and select a setting. Word wrap is supported.	Virtual Disks ↓ Virtual Disk 1
Underlined letter in a menu heading on the menu bar	Indicates a shortcut you can use by pressing <Alt> <underlined letter>. To use this feature, the menu must be activated. Menu shortcuts are allowed but cannot be used when a menu is active.	<u>A</u> dapter
Underlined letter in a menu item	Indicates a shortcut you can use to expand a menu by pressing <Alt> <menu underlined letter>. Pressing <Alt> again closes the menu. Menu shortcuts are allowed but cannot be used when a menu is active.	<u>V</u> irtual Disk 1
< >	Key presses are enclosed in angle brackets.	<F1>, <Esc>, <Enter>
<Enter>	After you highlight a menu item, press <Enter> to select that item. This opens an options menu for the menu item. This applies to only certain menu items, such as Virtual Disk # . In a list of options for that item, such as the write policy for a virtual disk, highlight a setting, such as write-through, and press <Enter> to select it. In the right frame, you can press <Enter> to select Tree View or List View under the View Type heading.	Select Add New VD and press <Enter> to create a new virtual disk.
<Esc>	After you expand a pop-up window, press <Esc> to close the window. You can continue to press <Esc> to exit the BIOS Configuration Utility.	Press <Esc> to return to the VD Mgmt screen.

Table 5-1. Menu Navigation Keys (continued)

Notation	Meaning and Use	Example
<Tab>	Press <Tab> to move the cursor to the next control on a Dialog or page.	Press <Tab> to move the cursor to the next parameter you want to change.
<Shift><Tab>	Press <Shift><Tab> to move the cursor to the previous control on a dialog or page.	Press <Shift><Tab> to move the cursor from Virtual Disk to Disk Group # .
<Ctrl><N>	Press <Ctrl><N> to move to the next menu screen among the main menu screens: VD Mgmt , PD Mgmt , Ctrl Mgmt , and Foreign View . When you return to the original menu, the cursor will be on the same menu item it was on before you pressed <Ctrl><N>.	Press <Ctrl><N> on the VD Mgmt screen to move to the PD Mgmt screen.
<Ctrl><P>	Press <Ctrl><P> to move to the previous menu screen among the main menu screens: VD Mgmt , PD Mgmt , Ctrl Mgmt , and Foreign View . When you return to the previous screen, the cursor will be on the same menu item it was on before you pressed <Ctrl><P>.	Press <Ctrl><P> on the PD Mgmt screen to return to the VD Mgmt screen.
<F1>	Press <F1> to access Help information. The Help screens display a glossary of topics you can use to access information about navigation, RAID levels, and general topics.	<F1>
<F2>	Press <F2> to access the context menu, which displays the list of options.	<F2>
<F12>	Press <F12> to display a list of controllers.	<F12>
Spacebar	Press the spacebar to select an item, such as a virtual disk in the List View , select all the virtual disks (Mark All), or deselect all the virtual disks (Unmark All).	Press the spacebar to select each virtual disk on which you want to check data consistency.

 **NOTE:** You may need to press <F5> to refresh the information on the screen.

RAID Configuration Functions

 **NOTE:** Dell OpenManage Storage Management can perform all the same tasks as and more tasks than the BIOS Configuration Utility.


After you attach physical disks, use a configuration utility to prepare a virtual disk. Groups of the SAS and Serial ATA (SATA) physical disks are organized into virtual disks that must be able to support the RAID level that you select. If the operating system is not yet installed, use the BIOS Configuration Utility to perform this procedure.

 **NOTE:** The PERC 5 controllers support Dell-qualified SATA physical disks.

Use the configuration utilities to perform the following tasks:

- Access controllers, virtual disks, and physical disks individually.
- Select a host controller to work on.
- Create hot spare physical disks.
- Configure virtual disks.
- Initialize one or more virtual disks.
- Rebuild failed physical disks.

The following sections describe the menu options and provide detailed instructions used to perform the configuration tasks. They apply to the BIOS Configuration Utility and Dell OpenManage Storage Management. The following is a list of the procedures used to configure physical disks into virtual disks.

- 1 Designate hot spares (optional).
See "Managing Dedicated Hot Spares" on page 64 for more information.
- 2 Define virtual disks using the group of physical disks.
 **NOTE:** A disk group is deleted when the last virtual disk in the disk group is deleted.
- 3 Save the configuration information.
- 4 Initialize the virtual disks.

See "Setting Up Virtual Disks" on page 58 for the detailed configuration procedures.

BIOS Configuration Utility Menu Options

This section describes the options for the BIOS Configuration Utility for each of the major menus:

- **Virtual Disk Management (VD Mgmt)** menu
- **Physical Disk Management (PD Mgmt)** menu
- **Controller Management (Ctrl Mgmt)** menu
- **Foreign Configuration View** menu

Most menus consist of two frames:

- A left frame with the menu options
- A right frame with details of the items selected in the left frame

The following sections describe the menu and submenu options for each of the major menus.

Virtual Disk Management (VD Mgmt)

The **Virtual Disk Management** screen, **VD Mgmt**, is the first screen that displays when you access the BIOS Configuration Utility. In the **Tree View**, the left frame displays the menus for the virtual disk management, which are:

- **Controller #**
 - **Disk Group #**
 - **Virtual Disks** (displayed in descending numerical order)
 - **Physical Disks**
 - **Space Allocation** (virtual disk size and free space you can use to create a virtual disk)
 - **Hot Spares** (global and dedicated)

In the **Tree View**, the right frame displays detailed information for the selected controllers, disk groups, virtual disks, physical disks, space allocation, and hot spares, as shown in Table 5-2.

Table 5-2. Information on the Virtual Disk Management Screen

Menu Item Selected in Left Frame	Information That Displays in Right Frame
Controller #	<ul style="list-style-type: none">• Number of disk groups (DG)• Number of virtuals disks (VD)• Number of physical disks (PD)
Disk Group #	<ul style="list-style-type: none">• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares
Virtual Disks	<ul style="list-style-type: none">• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available in the virtual disk• Number of free segments• Number of dedicated hot spares
Virtual Disk #	<ul style="list-style-type: none">• RAID level (0, 1, 5, 10, or 50)• RAID status of the virtual disk (failed, degraded, or optimal)• Size of the virtual disk• Stripe element size• Operation currently in progress

Table 5-2. Information on the Virtual Disk Management Screen (continued)

Menu Item Selected in Left Frame	Information That Displays in Right Frame
Physical Disks	<ul style="list-style-type: none">• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares
Physical Disk #	<ul style="list-style-type: none">• Vendor name• Physical disk size• Physical disk state• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares
Space Allocation	<ul style="list-style-type: none">• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares
Hot Spares	<ul style="list-style-type: none">• Vendor name• Physical disk size• Physical disk state• Number of virtuals disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares

In the **List View**, the **Virtual Disk Management** screen appears. In this view, the left frame displays the virtual disk number, RAID level, virtual disk status, virtual disk size, and stripe element size.

Use this screen to display the **Virtual Disk Parameters** window (used to select parameters, such as stripe element size and cache policy), and perform a consistency check on an individual virtual disk or multiple virtual disks. See "Setting Up Virtual Disks" on page 58 for the procedures used to perform these actions.

Virtual Disk Actions

Table 5-3 describes the actions you can perform on virtual disks. See "Setting Up Virtual Disks" on page 58 for procedures you can use to perform these actions.

Table 5-3. Virtual Disk Menu Options

Option	Description
Create a new virtual disk	Creates a new virtual disk from one or more physical disks.
Manage dedicated hot spares	Creates or deletes a hot spare that can be dedicated to a single redundant virtual disks.
Initialize a virtual disk	Initializes the selected virtual disk. Every virtual disk that is configured must be initialized.
Check data consistency on a virtual disk	Verifies the correctness of the redundancy data in the selected virtual disk. This option is available only if RAID level 1, 5, 10, or 50 is used. The PERC 5 controllers automatically correct any differences found in the data.
Display or update virtual disk parameters	Displays the properties of the selected virtual disk. You can modify the cache write policy, read policy, and the input/output (I/O) policy from this menu.
Delete a virtual disk	Deletes the virtual disk and frees up disk space to create another virtual disk.
Delete a disk group	Deletes a disk group, which is a collection of disks from one or more disk subsystems controlled by management software.

Physical Disk Management (PD Mgmt)

The **Physical Disk Management** screen, **PD Mgmt**, displays physical disk information and action menus. The screen displays physical disk IDs, vendor names, disk size, type, state, and disk group (**DG**). You can sort the list of physical disks based on these headings. You can perform several actions on the physical disks, including the following:

- Rebuilding physical disks
- Setting the LED to blink
- Making a disk online or unaffiliated
- Removing dedicated or global hot spares
- Creating global hot spares

Physical Disk Actions

Table 5-4 describes the actions you can perform on physical disks. See "Setting Up Virtual Disks" on page 58 for procedures that can be used to perform these actions.

Table 5-4. Physical Disk Options

Option	Description
Rebuild	Regenerates all data to a replacement disk in a redundant virtual disk (RAID level 1, 5, 10, or 50) after a disk failure. A disk rebuild normally occurs without interrupting normal operations on the affected virtual disk.
LED Blinking	Indicates when physical disks are being used to create a virtual disk. You can choose to start or stop the LED blinking.
Force Online	Changes the state of the selected physical disk to online.
Force Offline	Changes the state of the selected physical disk so that it is no longer part of a virtual disk.
Make Global HS	Designates the selected physical disk as a global hot spare. A global hot spare is part of a pool for all virtual disks controlled by the controller.
Remove HS	Removes a dedicated hot spare from its disk group or a global hot spare from the global pool of hot spares.

Rebuild

Select **Rebuild** to rebuild one or more failed physical disks. See "Performing a Manual Rebuild of an Individual Physical Disk" on page 67 for more information and the procedure to perform a physical disk rebuild.

Controller Management (Ctrl Mgmt)

The **Controller Management** screen, **Ctrl Mgmt**, displays the BIOS version, firmware version, and Ctrl-R version, and offers actions you can perform on the controller and BIOS. You can use this screen to enable or disable the controller BIOS, alarm, and the BIOS during bootup in event of BIOS errors. In addition, you can select a virtual disk from which to boot, select default settings, and reset the configuration.

Controller Management Disk Actions

Table 5-5 describes the actions you can perform on the **Ctrl Mgmt** screen. See "Setting Up Virtual Disks" on page 58 for procedures you can use to perform these actions.

Table 5-5. Controller Management Options

Option	Description
Enable Controller BIOS	Select this option to enable the controller BIOS. If the boot device is on the RAID controller, the BIOS must be enabled. Disable the BIOS to use other boot devices. In a multiple controller environment, you can enable BIOS on multiple controllers. However, if you want to boot from a specific controller, then enable the BIOS on that controller and disable it on the other controllers. The system can then boot from the BIOS-enabled controller.
Enable Alarm	Select this option to enable or disable the on-board alarm tone generator. The alarm sounds to alert you of key critical and warning events involving the virtual disk or physical disk problems. See Table 6-8 for details about the alarm.
Enable BIOS Stop On Error	Select this option to stop the system BIOS during bootup if there are BIOS errors. This option enables you to enter the configuration utility to resolve the problem. NOTE: Certain errors continue for five seconds even when Stop On Error is enabled.
Select Bootable Virtual Disk	Select this option to specify a virtual disk as the boot disk on the controller. This option displays if you have built virtual disks.
Select Defaults	Select this option to restore the default settings for the options in the Settings box.

Foreign Configuration View

Select **Foreign Configuration View** to display a foreign configuration before you either import the foreign configuration into the virtual disk or clear the foreign configuration. The section "Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen" on page 62 contains the procedures you can use to manage the foreign configurations.


Setting Up Virtual Disks

This section contains the procedures used to set up a disk group and create virtual disks. Each of the following procedures are explained individually in this section in detail.


- 1 Create the virtual disks and select the virtual disk options.

- 2 Designate hot spares (optional).

See "Managing Dedicated Hot Spares" on page 64 for more information.

 **NOTE:** A disk group is deleted when the last virtual disk in the disk group is deleted.

- 3 Initialize the virtual disks.

 **NOTE:** When you use one physical disk group to create multiple virtual disks, the virtual disks must all have the same RAID level.

When you define the virtual disks, you can set the virtual disk parameters described in Table 5-6.

- Stripe element size
- Write policy
- Read policy

Table 5-6. Virtual Disk Parameters and Descriptions

Parameter	Description
Stripe Element Size	<p>Stripe Element Size specifies the size of the segments written to each physical disk in a RAID 0, 1, 5, 10, and 50 virtual disk. You can set the stripe element size to 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB. The default and recommended stripe element size is 128 KB.</p> <p>A larger stripe element size provides better read performance, especially if your system does mostly sequential reads. However, if you are sure that your system does random read requests with small file sizes, select a small stripe element size.</p>
Write Policy	<p>Write Policy specifies the controller write policy. You can set the write policy to Write-Back or Write-Through.</p> <p>In Write-Back caching, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction.</p> <p>NOTE: The default cache setting is Write-back cache.</p> <p>NOTICE: If Write-back is enabled and the system is quickly turned off and then on, the controller may pause as the system flushes cache memory. Controllers that contain a battery backup will default to write-back caching.</p> <p>In Write-Through caching, the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction.</p> <p>NOTE: Certain data patterns and configurations perform better with Write-through cache policy.</p> <p>Write-Back caching has a performance advantage over write-through caching.</p>
Read Policy	<p>Read-ahead enables the read-ahead feature for the virtual disk. You can set this parameter to Read-ahead, No-read-ahead, or Adaptive. The default is No-read-ahead.</p> <p>Read-ahead specifies that the controller uses read-ahead for the current virtual disk. Read-ahead capability allows the controller to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the data will be needed soon. Read-ahead supplies sequential data faster, but is not as effective when accessing random data.</p> <p>No-read-ahead specifies that the controller does not use read-ahead for the current virtual disk.</p> <p>NOTE: No-read-ahead shows higher performance results due to the effectiveness of hard-drive caching algorithms.</p> <p>Adaptive specifies that the controller begins using read-ahead if the two most recent disk accesses occurred in sequential sectors. If all read requests are random, the algorithm reverts to No-read-ahead; however, all requests are still evaluated for possible sequential operation.</p>

Creating Virtual Disks

 **NOTE:** PERC 5 does not support creation of a virtual disk that combines SAS and SATA physical disks

Perform the following steps to create virtual disks.


- 1 During host system bootup, press <Ctrl><R> when the BIOS banner displays.
The **Virtual Disk Management** screen displays.

 **NOTE:** This procedure describes the BIOS Configuration Utility screens in **Tree View**.

- 2 Highlight **Controller #**.
- 3 Press <F2> to display the actions you can perform, which are creating new virtual disks, resetting a configuration, and importing or clearing a foreign configuration.
- 4 Select **Create New VD** and press <Enter>.

The **Create New VD** screen displays. The cursor is on the **RAID Levels** option.

- 5 Press <Enter> to display the RAID levels.
- 6 Press the down arrow key to select a RAID level and press <Enter>.
- 7 Press <Tab> to move the cursor to the list of physical disks.
- 8 Use the down arrow key to highlight a physical disk and press the spacebar to select the disk.
- 9 Select additional disks, if desired.
- 10 Press <Tab> to move the cursor to the box **Basic Settings** and set the virtual disk size and name.
The virtual disk size displays in megabyte (MB) format.



 **NOTE:** For RAID levels 0, 1, and 5 only, you can use part of the available disk space to create one virtual disk and then use the rest of the disk space to create another virtual disk or disks.

- 11 Press <Tab> to move the cursor to **Advanced Settings**.
- 12 Press the spacebar to make the settings active so that you can change them.

An X displays next to **Advanced Settings**. The settings are the stripe element size and read policy. You can also choose advanced options such as forcing the cache policy to be Write-back, initializing the virtual disk, and configuring a dedicated hot spare.

The defaults for these parameters display when the window displays. You can accept the defaults or change them. See "Virtual Disk Parameters and Descriptions" on page 59 for detailed information about the virtual disk parameters.


- 13 Perform the following steps to select the virtual disk parameters:
 - a Press <Tab> to move the cursor to the parameters you want to change.
 - b Press the down arrow key to open the parameters and scroll down the list of settings.
 - c To change the stripe element size, press <Tab> to highlight **Stripe Element Size**.
 - d Press <Enter> to display the list of stripe element sizes (8 KB, 16 KB, 32 KB, 64 KB, and 128 KB), then press the down arrow key to highlight an option and press <Enter>.

- e Press <Tab> to move the cursor to the **Read Policy** to change it if desired.
- f Press <Enter> to display the options, **Read Ahd**, **No Read**, or **Adaptive**, then press the down arrow key to highlight an option and press <Enter>.
- g Press <Tab> to move the cursor to the **Write Policy** to change it if desired.
- h Press <Enter> to display the options, **Write-Through** or **Write-Back**, then press the down arrow key to highlight an option and press <Enter>.
- i Press <Tab> to move the cursor to **OK**.
- j Press <Tab> to move the cursor to **Force WB with no battery** and press <Enter>.
- k Press <Tab> to move the cursor to **Initialize** and press <Enter>.
 -  **NOTE:** The initialization performed at this stage is fast initialization.
- l Press <Tab> to move the cursor to **Configure Hot Spare** and press <Enter>.
 -  **NOTE:** The hot spare created at this stage is a dedicated hot spare.
- m If you have chosen to create hot spares in the earlier steps a pop-up window appears where drives with appropriate sizes are displayed. Press the spacebar to select the drive size.
- n After you select the drive size, click **OK** to finalize the selection or click **Cancel** to forfeit the selection.
- o Select **OK** to accept the settings and press <Enter> to exit this window or select **Cancel** and press <Enter> to exit if you do not want to change any virtual disk parameters.

Initializing Virtual Disks

Perform the following steps to initialize virtual disks.


- 1 On the **VD Mgmt** screen, select **Virtual Disk #** and press <F2> to display the menu of available actions.
 - 2 Select **Initialization** and press the right arrow key to display the **Initialization** submenu options.
 - 3 Select **Start Init.** to begin a regular initialization or select **Fast Init.** to begin a fast initialization. Fast initialization writes zeroes to the first sector of the virtual disk so that initialization occurs in 2–3 seconds. A progress bar displays next to the selected virtual disk as it initializes.
 - 4 When initialization is complete, press <Esc> to return to the main menu screen.
 - 5 Repeat the procedures in this section to configure another virtual disk.
- The PERC 5 controllers support up to 64 virtual disks per controller. The currently configured virtual disks display on the screen.

 **NOTE:** Performing a full initialization eliminates the need for the virtual disk to undergo a background initialization. The full initialization is faster than the background initialization. Full initialization can be performed to speed up performance access to the virtual disk after initial creation. The full initialization must be completed before accessing the virtual disk.

Importing or Clearing Foreign Configurations Using the VD Mgmt Menu


When a foreign configuration exists, the BIOS banner displays the message `Foreign configuration(s) found on adapter`. In addition, a foreign configuration, when present, appears on the right side of the **VD Mgmt** screen.

You can use the **VD Mgmt** menu to import the existing configuration to the RAID controller or clear the existing configuration to create a new one. The information for the foreign configuration does not display until it is imported.

 **NOTE:** The controller does not allow an import that results in more than 64 virtual disks.

Perform the following steps to import or clear foreign configurations.

- 1 During bootup, press <Ctrl><R> when prompted by the BIOS banner.
The **VD Mgmt** screen appears by default.
- 2 On the **VD Mgmt** screen, highlight the **Controller #**.
The controller number is the only item that appears until you import the foreign configuration.
- 3 Press <F2> to display the available actions.
- 4 Press the right arrow key to display the available actions, **Import** and **Clear**.


 **NOTE:** Ensure that your virtual disk has all the physical disks by verifying that there are no physical disks marked as **Missing** in the foreign view page and that all the disks appear as expected before importing them.

- 5 Select **Import** to import the foreign configuration or **Clear** to delete the foreign configuration and then press <Enter>.

If you import the configuration, the **VD Mgmt** displays detailed configuration information. This includes information about the disk groups, virtual disks, physical disks, space allocation, and hot spares.

Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen

If one or more physical disks are removed from a configuration, for example, by a cable pull or physical disk removal, the configuration on those disks is considered a foreign configuration by the RAID controller. You can use the **Foreign Configuration View** screen to view detailed configuration information about the disk groups, virtual disks, physical disks, space allocation, and hot spares. After you view the foreign configuration, you can either import it to the RAID controller or clear it.

 **NOTE:** Before you import, review the configuration on the screen to ensure that it is the desired end result.

You can use the **Foreign Configuration View** screen to manage foreign configurations in the following cases:

- All the physical disks in a configuration are removed and re-inserted.
- Some of the physical disks in a configuration are removed and re-inserted.
- All the physical disks in a virtual disk are removed, but at different times, and then re-inserted.
- The physical disks in a non-redundant virtual disk are removed.

Perform the following procedures on the **Foreign Configuration View** screen to manage foreign configurations in each specific case:

- 1 If all or some of the physical disks in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations. Perform the following steps:
 - a Select **Foreign Configuration View** to display the foreign configuration information on the **Foreign Configuration View** screen.
 - b Press <F2> to display the options **Import** or **Clear**.



NOTE: You must have all the drives in the enclosure before you perform the import operation.

- c Select **Import** to import the foreign configuration to the controller or select **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).

Automatic rebuild occurs in redundant virtual disks if the VD transitions into DEGRADED before going OFFLINE.



NOTE: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. See "Checking Data Consistency" on page 66 for more information about checking data consistency.

- 2 If all the physical disks in a virtual disk are removed, but at different times, and re-inserted, the controller considers the drives to have foreign configurations. Perform the following steps:

- a Select **Foreign Configuration View** to display the complete virtual disk, across different foreign configurations and allow foreign configurations to be imported.
 - b Press <F2> to display the options **Import** or **Clear**.



NOTE: You must have all the drives in the enclosure before you perform the import operation.

- c Select **Import** to merge the foreign configurations with the existing configuration on the controller or **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).

If you select **Import**, all drives that were pulled before the virtual disk became offline will be imported, and then automatically rebuilt. Automatic rebuild occurs in redundant virtual disks if the VD transitions into DEGRADED before going OFFLINE.



NOTE: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. See "Checking Data Consistency" on page 66 for more information about checking data consistency.

- 3 If the physical disks in a non-redundant virtual disk are removed, the controller considers the drives to have foreign configurations. Perform the following steps:

- a Select **Foreign Configuration View** to display the complete foreign configuration information.
 - b Press <F2> to display the options **Import** or **Clear**.
 - c Select **Import** to import the foreign configuration to the virtual disk or **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).

No rebuilds will occur after the import operation because there is no redundant data to rebuild the drives with.

Setting LED Blinking

The LED blinking option indicates when physical disks are being used to create a virtual disk. You can choose to start or stop the LED blinking. Perform the following steps to start or stop this option.

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.
A list of physical disks appears. The status of the each disk displays under the heading **State**.
- 2 Press the down arrow key to highlight a physical disk.
- 3 Press <F2> to display the menu of available actions.
- 4 Press the down arrow key to highlight **LED Blinking**.
- 5 Press the right arrow key to display the available actions, **Start** and **Stop**.
- 6 Select **Start** to begin LED blinking or **Stop** to end LED blinking.


Managing Dedicated Hot Spares

A dedicated hot spare automatically replaces a failed physical disk only in the selected disk group which the hot spare is part of. A dedicated hot spare is used before one of the global hot spares is used. You can create dedicated hot spares or delete them on the **VD Mgmt** screen. Perform the following steps to create or delete dedicated hot spares.

- 1 On the **VD Mgmt** screen, select **Disk Group #** and press <F2> to display the menu of available actions.
The available menu options appear.


- 2 Select **Manage Ded. HS** and press <Enter>.

A screen displays a list of the current dedicated hot spares and the physical disks that are available to create dedicated hot spares. An X displays next to the current dedicated hot spares.

 **NOTE:** The utility allows only disks of the same drive technology and of equal or greater size to be selected as dedicated hot spare.


- 3 To create a dedicated hot spare, press the down arrow key to highlight an available physical disk and press the spacebar to select the disk. Do this for as many dedicated hot spares as you want to create.
An X displays next to the selected physical disk(s).
- 4 To delete a dedicated hot spare, use the down arrow key to highlight a current hot spare and press the spacebar to deselect the disk. Do this for as many dedicated hot spares as you want to delete.
- 5 Press <Enter> to approve the changes.

The **VD Mgmt** screen displays the updated list of hot spares under the **Hot spares** heading.

 **NOTE:** If a dedicated hot spare is removed, re-inserted, and then imported, the array disk transitions to a global hot spare upon completion of the foreign configuration import.



Creating Global Hot Spares

A global hot spare can be used to replace a failed physical disk in any redundant array as long as the capacity of the global hot spare is equal to or larger than the coerced capacity of the failed physical disk. Perform the following steps to create global hot spares.

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.
A list of physical disks displays. The status of the each disk displays under the heading **State**.
- 2 Press the down arrow key to highlight a physical disk to change to a global hot spare.
- 3 Press <F2> to display the menu of available actions.
- 4 Press the down arrow key to highlight **Make Global HS** and press <Enter>.
The physical disk is changed to a global hot spare. The status of the physical disk as a global hot spare displays under the heading **State**.
 **NOTE:** To replace a failed physical disk global hot spares must use the same drive technology and must be equal or greater in size.
- 5 Select additional physical disks if desired and follow the previous steps to change them to global hot spares.

Removing Global or Dedicated Hot Spares

You can remove one global or dedicated hot spare at a time on the **PD Mgmt** screen. Perform the following steps to remove a global hot spare or dedicated hot spare.

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.
A list of physical disks displays. The status of each disk displays under the heading **State**.
- 2 Press the down arrow key to highlight a physical disk that is a hot spare.
- 3 Press <F2> to display the menu of available actions.
- 4 Press the down arrow key to select **Remove HS** from the list of actions and press <Enter>.
The physical disk is changed to the **Ready** state. The status of the physical disk displays under the heading **State**.
 **NOTE:** Try to use physical disks of the same capacity in a specific virtual disk. If you use physical disks with different capacities in a virtual disk, all physical disks in the virtual disk are treated as if they have the capacity of the smallest physical disk.
- 5 Select additional hot spares if desired and follow the previous steps to remove them.
 **NOTE:** On the PERC 5/i RAID controller, after you delete a virtual disk, any hot spares dedicated to that virtual disk become global hot spares.

Enabling the Alarm to Alert in Case of Physical Disk Failures

The PERC 5/E Adapter uses an audible alarm to alert you to key critical and warning events involving the virtual disks or physical disks. You can use the BIOS Configuration Utility to enable or disable the on-board alarm tone. See Table 6-8 for the list of critical and warning events, severity levels of the events, and audible codes.

Perform the following steps to enable the alarm in the BIOS Configuration Utility.

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> in the **Settings** box to move the cursor to **Enable Alarm**.
- 3 Press the spacebar to select **Enable Alarm**.
An **X** displays next to **Enable Alarm**.
- 4 Press <Tab> to move the cursor to the **Apply** button, then press <Enter> to apply the selection.
The alarm is enabled. To disable the alarm, use the spacebar to deselect the **Enable Alarm** control, and then select **Apply**.

Checking Data Consistency

Select the **Consistency Check** option in the configuration utility to verify the redundancy data in virtual disks that use RAID levels 1, 5, 10, and 50. (RAID 0 does not provide data redundancy.)



NOTE: Dell recommends that you run data consistency checks on a redundant array at least once a month. This allows detection and automatic replacement of bad blocks. Finding a bad block during a rebuild of a failed physical disk is a serious problem, as the system does not have the redundancy to recover the data.



NOTE: The system will take longer to reboot after you perform a data consistency check.

Perform the following steps to run a **Consistency Check**.

- 1 Press <Ctrl><N> to access the **VD Mgmt** menu screen.
- 2 Press the down arrow key to highlight **Virtual Disk #**.
- 3 Press <F2> to display the menu of available actions.
- 4 Press the down arrow key to select **Consistency Check**.
- 5 Press the right arrow key to display the menu of available actions, **Start** and **Stop**.
- 6 Select **Start** and press <Enter> to run a **Consistency Check**.
The **Consistency Check** runs and checks the redundancy data in the virtual disks.
- 7 After you start the **Consistency Check**, press <Esc> to display the previous menu.

Stopping Background Initialization

Background initialization is the automated check for media errors in which parity is created and written in the process. BGI does not run on RAID 0 virtual disks. Under certain conditions, Ctrl-R will ask you if you want to stop BGI in progress. An alert message displays if BGI is in progress and you start any of the following actions:

- A full initialization on the virtual disk
- A quick initialization on the virtual disk
- A consistency check on the virtual disk

The following alert message displays: The virtual disk is undergoing a background initialization process. Would you like to stop the operation and proceed with the <full initialization/quick initialization/consistency check> instead?

Click **Yes** to stop the BGI and start the requested operation or **No** to allow the BGI to continue.

Performing a Manual Rebuild of an Individual Physical Disk

Use the following procedures to rebuild one failed physical disk manually.

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.

A list of physical disks appears. The status of each disk appears under the heading **State**.

- 2 Press the down arrow key to highlight a physical disk that has a failed state.

- 3 Press <F2> to display a menu of available actions.

The **Rebuild** option is highlighted at the top of the menu.

- 4 Press the right arrow key to display the rebuild options and select **Start**.

- 5 After you start the rebuild, press <Esc> to display the previous menu.



NOTICE: If a rebuild operation is in progress on a physical disk that is a member of a disk group that contains multiple virtual disks, such as a sliced configuration, and one of the virtual disks is deleted while the rebuild is in progress, then the rebuild operation will stop. If this occurs, resume the rebuild operation manually using a storage management application. In order to avoid interruption, ensure that none of the virtual disks that reside on the rebuilding physical disk are deleted until the rebuild operation is complete.

Deleting Virtual Disks

To delete virtual disks, perform the following steps in the BIOS Configuration Utility.



NOTE: You will not be allowed to delete a virtual disk during an initialization.



NOTE: Warning messages display the effect of deleting a virtual disk. You must accept two warning statements before the virtual disk deletion is completed.

- 1 Press <Ctrl><N> to access the **VD Mgmt** screen.


- 2 Press <Tab> to move the cursor to a virtual disk under the **Virtual Disks** heading.

- 3 Press <F2>.

The action menu appears.

- 4 Select **Delete VD** and press <Enter>.

- 5 On the **VD Mgmt** screen, select **Space Allocation** to display the amount of free space available after you delete the virtual disk.

 **NOTE:** On the PERC 5/i RAID controller, after you delete a virtual disk, any hot spares dedicated to that virtual disk become global hot spares.

Deleting Disk Groups

You can delete disk groups using the BIOS Configuration Utility. When you delete a disk group, the utility also removes the virtual disks in that disk group.

To delete disk groups, perform the following steps in the BIOS Configuration Utility.

- 1 Press <Ctrl><N> to access the **VD Mgmt** screen.
- 2 Press <Tab> to move the cursor to a disk group under the **Virtual Disks** heading.
- 3 Press <F2>.

The action menu displays.

- 4 Select **Delete Disk Group** and press <Enter>.


This deletes the disk group. When you delete a disk group, the remaining disk groups with higher numbers are automatically renumbered. For example, if you delete disk group #2, disk group #3 is then automatically renumbered as disk group #2.

Upgrading Firmware

You can download the latest firmware from the Dell Support website and flash it to the firmware on the controller.

 **NOTE:** Go to the Dell Support website at support.dell.com.

Enabling Boot Support

 **NOTE:** See your system documentation to ensure the proper boot order is selected in the system BIOS.

In a multiple controller environment, multiple controllers can have their BIOS enabled. However, if you want to boot from a specific controller, then enable the BIOS on that controller and disable it on the other controllers. The system can then boot from the BIOS-enabled controller. Perform the following steps to enable the controller BIOS.

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> to move the cursor to **Enable Controller BIOS** in the **Settings** box.

- 3 Press the spacebar to select **Enable Controller BIOS**.
An **X** displays next to **Enable Controller BIOS**.
- 4 Press <Tab> to move the cursor to the **Apply** button, and then press <Enter> to apply the selection.
The controller BIOS is enabled. To disable the controller BIOS, use the spacebar to deselect the **Enable Controller BIOS** control, and then select **Apply** and press <Enter>.


After you enable the BIOS for a controller, perform the following steps to enable the boot support for that controller.

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> to move the cursor to the **Select Bootable VD** in the **Settings** box.
Shadow lines display if there are two or more virtual disks.
- 3 Press the down arrow key to display a list of virtual disks.
- 4 Use the down arrow key to move the cursor to a virtual disk.
- 5 Press <Enter> to select a virtual disk.
- 6 Press <Tab> to move the cursor to the **Apply** button, and then press <Enter> to apply the selection.
Boot support is enabled for the selected controller.

Enabling BIOS Stop on Error

The option **BIOS Stop on Error** is used to stop the system from booting if there are BIOS errors. Perform the following steps to enable **BIOS Stop on Error**.

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> to move the cursor to **Enable BIOS Stop on Error** in the **Settings** box.
- 3 Press the spacebar to select **Enable BIOS Stop on Error**.
An **X** displays next to **Enable BIOS Stop on Error**.
- 4 Press <Tab> to move the cursor to the **Apply** button, and then press <Enter> to apply the selection.
The controller BIOS is enabled. To disable **Enable BIOS Stop on Error**, use the spacebar to deselect **Enable BIOS Stop on Error**, then select **Apply** and press <Enter>.

 **NOTE:** Certain errors continue for five seconds even when **Stop on Error** is enabled.

Restoring Factory Default Settings

You can use the **Ctrl Mgmt** menu screen to restore the default settings for the options in the **Settings** box. The settings are **Enable Controller BIOS**, **Enable Alarm**, and **Enable BIOS Stop on Error**. Perform the following steps to restore default settings.

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> to move the cursor to the **Settings** box.

- 3** Use the spacebar to deselect the settings for the options in the **Settings** box.
- 4** Press <Tab> to move the cursor to the **Advanced** box.
- 5** Press the spacebar to select **Advanced**.
An **X** displays next to **Advanced**.
- 6** Press <Tab> to move the cursor to **Select Defaults**.
- 7** Press the spacebar to select **Select Defaults**.
The defaults are automatically selected for the controller settings and display in the **Settings** box.
An **X** displays by each setting which is enabled by default.

Troubleshooting

To get help with your Dell™ PowerEdge™ Expandable RAID Controller (PERC) 5 controller, you can contact your Dell Technical Service representative or access the Dell Support website at support.dell.com.

Virtual Disks Degraded

A redundant virtual disk is in a degraded state when one physical disk has failed or is inaccessible. For example, a RAID 1 virtual disk consisting of two physical disks can sustain one physical disk in a failed or inaccessible state and become a degraded virtual disk.

To recover from a degraded virtual disk, rebuild the physical disk in the inaccessible state. Upon successful completion of the rebuild process, the virtual disk state changes from degraded to optimal. For the rebuild procedure, see "Performing a Manual Rebuild of an Individual Physical Disk" on page 67 in "RAID Configuration and Management" on page 49.

Memory Errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from these memory errors. Single-bit memory errors can be handled by the firmware and do not disrupt normal operation. A notification will be sent if the number of single-bit errors exceeds a threshold value.

Multi-bit errors are more serious, as they result in corrupted data and data loss. The following are the actions that occur in the case of multi-bit errors:

- If an access to data in cache memory causes a multi-bit error when the controller is started with dirty cache, the firmware will discard the cache contents. The firmware will generate a warning message to the system console to indicate that the cache was discarded and will generate an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the firmware will stop.
- The firmware will log an event to the firmware internal event log and will log a message during POST indicating that a multi-bit error has occurred.



NOTE: In case of a multi-bit error, contact Dell Technical Support.

General Problems

Table 6-1 describes general problems you might encounter, along with suggested solutions.

Table 6-1. General Problems

Problem	Suggested Solution
The device displays in Device Manager but has a yellow bang (exclamation point).	Reinstall the driver. See the driver installation procedures in the section "Driver Installation" on page 39.
The device does not appear in Device Manager.	Turn off the system and reseal the controller.
No Hard Drives Found message appears during a CD installation of Microsoft® Windows® 2000 Server, Windows Server® 2003, or Windows XP because of the following causes: 1 The driver is not native in the operating system. 2 The virtual disks are not configured properly. 3 The controller BIOS is disabled.	The corresponding solutions to the three causes of the message are: 1 Press <F6> to install the RAID Device Driver during installation. 2 Enter the BIOS Configuration Utility to configure the virtual disks. See the section "RAID Configuration and Management" on page 49 for procedures to configure the virtual disks. 3 Enter the BIOS Configuration Utility to enable the BIOS. See the section "Hardware Installation and Configuration" on page 25 for procedures to configure the virtual disks.

Physical Disk Related Issues

Table 6-2 describes physical disk-related problems you might encounter, along with suggested solutions.

Table 6-2. Physical Disk Issues

Problem	Suggested Solution
One of the physical disks in the disk array is in the inaccessible state.	Perform the following actions to resolve this problem: <ul style="list-style-type: none">• Check the enclosure or backplane for damage.• Check the SAS cables.• Reseat the physical disk.• Contact Dell Technical Support if the problem persists.
Cannot rebuild a fault tolerant virtual disk.	This could result from any of the following: <ul style="list-style-type: none">• The replacement disk is too small. Replace the failed disk with a good physical disk with sufficient capacity.
Fatal errors or data corruption are reported when accessing virtual disks.	Contact Dell Technical Support.

Physical Disk Failures and Rebuilds

Table 6-3 describes issues related to physical disk failures and rebuilds.

Table 6-3. Physical Disk Failure and Rebuild Issues

Issue	Suggested Solution
Rebuilding a physical disk after one of them is in an inaccessible state.	<p>If you have configured hot spares, the PERC 5 controller automatically tries to use one to rebuild a physical disk that is in an inaccessible state. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the inaccessible physical disks are available. You must insert a physical disk with enough storage into the subsystem before rebuilding the physical disk. You can use the BIOS Configuration Utility or Dell OpenManage™ Storage Management application to perform a manual rebuild of an individual physical disk.</p> <p>See the section "Performing a Manual Rebuild of an Individual Physical Disk" on page 67 in "RAID Configuration and Management" on page 49 for procedures to rebuild a single physical disk.</p>
Rebuilding the physical disks after multiple disks become simultaneously inaccessible.	<p>Multiple physical disk errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data. It is possible to recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk.</p> <ol style="list-style-type: none"> 1 Turn off the system, check cable connections, and reseal physical disks. <ul style="list-style-type: none"> Follow the safety precautions to prevent electrostatic discharge. Ensure that all the drives are present in the enclosure. 2 Power up the system and enter into the CTRL-R utility and import the foreign configuration. <ul style="list-style-type: none"> If the VD is redundant and transitioned into DEGRADED state before going OFFLINE a rebuild operation starts automatically after the configuration is imported. If the VD has gone directly into the OFFLINE state due to a cable pull or power loss situation the VD will be imported in its OPTIMAL state without a rebuild occurring. <p>You can use the BIOS Configuration Utility or Dell OpenManage Storage Management application to perform a manual rebuild of multiple physical disks.</p> <p>See the section "Performing a Manual Rebuild of an Individual Physical Disk" on page 67 in "RAID Configuration and Management" on page 49 for procedures to rebuild a single physical disk.</p>
A virtual disk fails during rebuild while using a global hot spare.	The global hot spare goes back into HOTSPARE state and the virtual disk goes into FAIL state.
A virtual disk fails during rebuild while using a dedicated hot spare.	The dedicated hot spare goes into READY state and the virtual disk goes into FAIL state.

Table 6-3. Physical Disk Failure and Rebuild Issues (continued)

Issue	Suggested Solution
A physical disk becomes inaccessible during a reconstruction process on a redundant virtual disk that has a hot spare.	The rebuild operation for the inaccessible physical disk starts automatically after the reconstruction is completed.
A physical disk is taking longer than expected to rebuild.	A physical disk takes longer to rebuild when under high stress. For example, there is one rebuild input/output (I/O) operation for every five host I/O operations.

SMART Error

Table 6-4 describes issues related to the Self-Monitoring Analysis and Reporting Technology (SMART). SMART monitors the internal performance of all motors, heads, and physical disk electronics and detects predictable physical disk failures.



NOTE: For information about where to find reports of SMART errors that could indicate hardware failure, see the Dell OpenManage Storage Management documentation.

Table 6-4. SMART Error

Problem	Suggested Solution
A SMART error is detected on a physical disk in a redundant virtual disk.	Perform the following steps: <ol style="list-style-type: none">1 Force the physical disk offline.2 Replace it with a new physical disk of equal or higher capacity.3 Perform a rebuild. See "Performing a Manual Rebuild of an Individual Physical Disk" on page 67 for rebuild procedures.
A SMART error is detected on a physical disk in a non-redundant virtual disk.	Perform the following steps: <ol style="list-style-type: none">1 Back up your data.2 Delete the virtual disk. See "Deleting Virtual Disks" on page 67 for information on deleting a virtual disk.3 Replace the affected physical disk with a new physical disk of equal or higher capacity.4 Recreate the virtual disk. See "Setting Up Virtual Disks" on page 58 for information on creating virtual disks.5 Restore the backup.

PERC 5 Post Error Messages

In PERC 5 controllers, the BIOS (read-only memory, ROM) provides INT 13h functionality (disk I/O) for the virtual disks connected to the controller, so that you can boot from or access the physical disks without the need of a driver. Table 6-5 describes the error messages and warnings that display for the BIOS.

Table 6-5. BIOS Errors and Warnings

Message	Meaning
BIOS Disabled. No Logical Drives Handled by BIOS	This warning displays after you disable the ROM option in the configuration utility. When the ROM option is disabled, the BIOS cannot hook Int13h and cannot provide the ability to boot from the virtual disk. (Int13h is an interrupt signal that supports numerous commands that are sent to the BIOS, then passed to the physical disk. The commands include actions you can perform with a physical disk, such as reading, writing, and formatting.)
Press <Ctrl><R> to Enable BIOS	When the BIOS is disabled, you are given the option to enable it by entering the configuration utility. You can change the setting to Enabled in the configuration utility.
Adapter at Baseport xxxx is not responding where xxxx is the baseport of the controller	If the controller does not respond for any reason but is detected by the BIOS, it displays this warning and continues. Shut down the system and try to reseat the controller. If this message still occurs, contact Dell Technical Support.
x Virtual Disk(s) Failed where x is the number of virtual disks failed	When the BIOS detects virtual disks in the failed state, it displays this warning. You should check to determine why the virtual disks failed and correct the problem. No action is taken by the BIOS.
x Virtual Disk(s) Degraded where x is the number of virtual disks degraded	When the BIOS detects virtual disks in a degraded state, it displays this warning. You should try to make the virtual disks optimal. No action is taken by the BIOS.
Memory/Battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue.	This message occurs under the following conditions: <ul style="list-style-type: none"> • The adapter detects that the cache in the controller cache has not yet been written to the disk subsystem • The controller detects an error-correcting code (ECC) error while performing its cache checking routine during initialization • The controller then discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed To resolve this problem, allow the battery to charge fully. If the problem persists, the battery or adapter DIMM might be faulty. In that case, contact Dell Technical Support.
Firmware is in Fault State	Contact Dell Technical Support.

Table 6-5. BIOS Errors and Warnings (continued)

Message	Meaning
Firmware version inconsistency was detected. The adapter has recovered, but cached data was lost. Press any key to continue.	New firmware has been flashed that is incompatible with the previous version. The cache contains data that has not been written to the physical disks and that cannot be recovered. Check data integrity. You may need to restore the data from a backup.
Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility.	When a controller firmware detects a physical disk with existing foreign metadata, it flags the physical disk as <i>foreign</i> and generates an alert indicating that a foreign disk was detected. You can use the BIOS Configuration Utility to import or clear the foreign configuration.
The foreign configuration message is always present during POST but no foreign configurations are present in the foreign view page in CTRL+R and all virtual disks are in an optimal state.	Clear the foreign configuration using CTRL+R or Dell OpenManage™ Server Administrator Storage Management. If a physical disk is inserted into the system that was once a member of a virtual disk, and that disk's previous location has been taken by a replacement disk through a rebuild, the newly inserted disk that was once a member of the virtual disk must have its foreign configuration flag manually removed.
Previous configuration(s) cleared or missing. Importing configuration created on XX/XX XX.XX. Press any key to continue, or 'C' to load the configuration utility.	The message means the controller and physical disks have different configurations. You can use the BIOS Configuration Utility to import or clear the foreign configuration.
There are X enclosures connected to port X but only X may be connected to a single SAS port. Please remove the extra enclosures then restart your system.	Too many enclosures are attached to one port. The extra enclosures must be removed and the system restarted.

Table 6-5. BIOS Errors and Warnings (continued)

Message	Meaning
Invalid SAS topology detected. Please check your cable configurations, repair the problem, and restart your system.	The SAS cables for your system are improperly connected. Check the cable connections and fix any problems, then restart the system. You may need to restore your data from a backup.
Multi-bit errors are detected on the controller. DIMM on the controller needs replacement. If you continue, data corruption can occur. Press 'X' to continue or else power off the system and replace the DIMM module and reboot. If you have replaced the DIMM please press 'X' to continue.	There are multi-bit ECC errors (MBE). ECC errors are errors that occur in the memory, which can corrupt cached data so that it has to be discarded. NOTICE: MBE errors are serious, as they result in corrupted data and data loss. In case of MBE errors, contact Dell Technical Support. NOTE: A similar message appears when multiple single-bit ECC errors are detected on the controller during bootup.
Some configured disks have been removed from your system, or are no longer accessible. Check your cables and ensure all disks are present. Press any key or 'C' to continue.	An array has failed. Some configured disks were removed from the system or, if not removed, are no longer accessible for other reasons. The SAS cables for your system might be improperly connected. Check the cable connections and fix any problems, then restart the system. You may need to restore your data from a backup. If there are no cable problems, press any key or <C> to continue.
Physical disk removed: Physical Disk {x.x.x} Controller {x}, Connector {x}	These two messages appear in the event log when you remove a drive. One indicates that the disk was removed and the other indicates that the device has failed. This is expected behavior.
Device failed: Physical Disk {x.x.x} Controller {x}, Connector {x}."	A storage component such as a physical disk or an enclosure has failed. The failed component might have been identified by the controller while performing a task such as a rescan or a check consistency. Replace the failed component. You can identify which disk has failed by locating the disk that has a red "X" for its status. Perform a rescan after replacing the disk.
Battery is missing or the battery could be fully discharged. If battery is connected and has been allowed to charge for 30 minutes and this message continues to appear, then contact Technical Support for assistance.	<ul style="list-style-type: none">• The controller battery is missing or damaged.• The controller battery is completely discharged and needs to be charged for it to become active. The battery must first be charged and the system must be restarted for the battery to be active again.

Red Hat Enterprise Linux Operating System Errors

Table 6-6 describes an issue related to the Red Hat® Enterprise Linux operating system.

Table 6-6. Linux Operating System Error

Error Message	Suggested Solution
<pre><Date:Time> <HostName> kernel: sdb: asking for cache data failed <Date:Time> <HostName> kernel: sdb: assuming drive cache: write through</pre>	<p>This error message displays when the Linux Small Computer System Interface (SCSI) mid layer asks for physical disk cache settings. Because the PERC 5 controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, the firmware does not respond to this command. Thus, the Linux SCSI mid layer assumes that the virtual disk's cache policy is write-through. SDB is the device node for a virtual disk. This value changes for each virtual disk.</p> <p>See the section "Setting Up Virtual Disks" on page 58 for more information about write-through cache.</p> <p>Except for this message, there is no side effect to this behavior. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC5 SAS RAID system remain the settings you have already chosen.</p>
<pre>Driver does not auto-build into new kernel after customer updates.</pre>	<p>This error is a generic problem for DKMS and applies to all DKMS-enabled driver packages. This issue occurs when you perform the following steps:</p> <ol style="list-style-type: none"> 1 Install a DKMS-enabled driver package. 2 Run <code>up2date</code> or a similar tool to upgrade the kernel into the latest version. 3 Reboot into the new kernel. <p>The driver running in the new kernel is the native driver in the new kernel. The driver package you once installed in the new kernel does not take effect in the new kernel.</p> <p>Perform the following procedure to make the driver auto-build into the new kernel:</p> <ol style="list-style-type: none"> 1 Type: <pre>dkms build -m <module_name> -v <module version> -k <kernel version></pre> 2 Type: <pre>dkms install -m <module_name> -v <module version> -k <kernel version></pre> 3 Type the following to check whether the driver is successfully installed in the new kernel: <pre>DKMS</pre> <p>The following details appear:</p> <pre><driver name>, <driver version>, <new kernel version>: installed</pre>

Table 6-6. Linux Operating System Error (continued)

Error Message	Suggested Solution
<pre>smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf</pre>	<p>These error messages are caused by an unsupported command coming directly from the user application. This is a known issue in which user applications try to direct Command Descriptor Blocks to RAID volumes. This error message has no effect on the user and there is no loss of functionality due to this error.</p> <p>The Mode Sense/Select command is supported by firmware on the PERC 5. However, the Linux kernel daemon is issuing the command to the virtual disk instead of to the driver IOCTL node. This action is not supported.</p>

LED Behavior Patterns

The external SAS ports on the PERC 5/E Adapter have a port status LED per x4 SAS port. This bi-color LED displays the status of any external SAS port. The LED indicates whether all links are functional or only partial links are functional. Table 6-7 describes the patterns for the port status.

Table 6-7. LED Behavior Patterns

Port State	LED State
Power-on state	Off
Reset state	Off
All links in port connected	Green light on
One or more links are not connected (applicable only in wide port configurations)	Amber light on
All links in the port are disconnected or the cable is disconnected	Off

Audible Alarm Warnings

An audible alarm is available on the PERC 5/E Adapter to alert you of key critical and warning events involving the virtual disk or physical disk problems. You can use the Basic Input/Output System (BIOS) Configuration Utility to enable, disable, or silence the on-board alarm tone.



NOTE: Silencing the alarm stops only the current alarm, but future alarms will be sounded. To permanently disable the alarm, select the **disable alarm** option.

Table 6-8 lists the critical and warning events, severity levels of the events, and audible codes.

Table 6-8. Audible Alarm Descriptions

Description	Severity	Audible Code
Controller alarm enabled	Normal	N/A
Virtual disk failed	Critical	3 seconds on, 1 second off
Virtual disk degraded	Warning	1 second on, 1 second off
Global hot spare failed	Warning	1 second on, 1 second off
Dedicated hot spare failed	Warning	1 second on, 1 second off
Physical disk failed	Critical	1 second on, 1 second off
Rebuild completed on physical disk	Normal	1 second on, 3 seconds off
Rebuild failed on physical disk	Warning	1 second on, 1 second off
Physical disk offline	Critical	1 second on, 1 second off



NOTE: If the PERC 5/E alarm was already beeping due to a previous failure and a new virtual disk is created on the same controller, then the previous alarm will be silenced. This is expected behavior.

Appendix: Regulatory Notices

Regulatory Notices

Electromagnetic Interference (EMI) is any signal or emission, radiated in free space or conducted along power or signal leads, that endangers the functioning of radio navigation or other safety service or seriously degrades, obstructs, or repeatedly interrupts a licensed radio communications service. Radio communications services include but are not limited to AM/FM commercial broadcast, television, cellular services, radar, air-traffic control, pager, and Personal Communication Services (PCS). These licensed services, along with unintentional radiators such as digital devices, including computer systems, contribute to the electromagnetic environment.

Electromagnetic Compatibility (EMC) is the ability of items of electronic equipment to function properly together in the electronic environment. While this computer system has been designed and determined to be compliant with regulatory agency limits for EMI, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause interference with radio communications services, which can be determined by turning the equipment off and on, you are encouraged to try to correct the interference by one or more of the following measures:

- Reorient the receiving antenna.
- Relocate the computer with respect to the receiver.
- Move the computer away from the receiver.
- Plug the computer into a different outlet so that the computer and the receiver are on different branch circuits.

If necessary, consult a Dell™ Technical Support representative or an experienced radio/television technician for additional suggestions.

Dell computer systems are designed, tested, and classified for their intended electromagnetic environment. These electromagnetic environment classifications generally refer to the following harmonized definitions:

- Class A is typically for business or industrial environments.
- Class B is typically for residential environments.

Information Technology Equipment (ITE), including peripherals, expansion cards, printers, input/output (I/O) devices, monitors, and so on, that are integrated into or connected to the system should match the electromagnetic environment classification of the computer system.

A Notice about Shielded Signal Cables: Use only shielded cables for connecting peripherals to any Dell device to reduce the possibility of interference with radio communications services. Using shielded cables ensures that you maintain the appropriate EMC classification for the intended environment. For parallel printers, a cable is available from Dell. If you prefer, you can order a cable from Dell on the World Wide Web at www.dell.com.

To determine the electromagnetic classification for your system or device, see the following sections specific for each regulatory agency. Each section provides country-specific EMC/EMI or product safety information.

This Regulatory appendix covers the following Dell products:

Please find a list of Class A EMC Environmental products:

- PERC 5/E Adapter
- PERC 5/i Integrated

Please find a list of Class B EMC Environmental products:

- PERC 5/i Adapter

FCC Notices (U.S. Only)

FCC, Class A

This product has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This product generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the manufacturer's instruction manual, may cause harmful interference with radio communications. Operation of this product in a residential area is likely to cause harmful interference, in which case you will be required to correct the interference at your own expense.

FCC, Class B

This product generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the manufacturer's instruction manual, may cause interference with radio and television reception. This product has been tested and found to comply with the limits for a Class B digital device pursuant to Part 15 of the FCC Rules.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

- 1 This device may not cause harmful interference.
- 2 This device must accept any interference received, including interference that may cause undesired operation.



NOTICE: The FCC regulations provide that changes or modifications not expressly approved by Dell Inc. could void your authority to operate this equipment.

These limits are designed to provide reasonable protection against harmful interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference with radio or television reception, which can be determined by turning the equipment off and on, you are encouraged to try to correct the interference by one or more of the following measures:

- Reorient the receiving antenna.
- Relocate the system with respect to the receiver.
- Move the system away from the receiver.
- Plug the system into a different outlet so that the system and the receiver are on different branch circuits.

If necessary, consult a representative of Dell Inc. or an experienced radio/television technician for additional suggestions.

The following information is provided on the device or devices covered in this document in compliance with FCC regulations:

Product Name:	Dell PERC 5/i Adapter
Company Name:	Dell Inc. Worldwide Regulatory Compliance & Environmental Affairs One Dell Way Round Rock, Texas 78682 USA 512-338-4400

Industry Canada (Canada Only)

Industry Canada, Class A


This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Industry Canada, Class B

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

 **NOTICE:** The Industry Canada regulations provide that changes or modifications not expressly approved by Dell Inc. could void your authority to operate this equipment.

CE Notice (European Union)

CE Notice (European Union)

This product has been determined to be in compliance with 73/23/EEC (Low Voltage Directive), 89/336/EEC (EMC Directive), and amendments of the European Union.

European Union, Class A

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

European Union, Class B

This Dell device is classified for use in a typical Class B domestic environment.

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

Předpisy CE (Evropská unie)

Bylo ověřeno, že tento produkt vyhovuje směrnicím 73/23/EEC (nízkonapěťová směrnice), 89/336/EEC (směrnice EMC) a dodatkům Evropské unie.

Evropská unie, třída A

RÁDIOVÉ RUŠENÍ - UPOZORNĚNÍ: Toto je produkt třídy A. V domácnosti toto zařízení může způsobovat rádiové rušení (RF). V tom případě bude nutné, aby uživatel podnikl příslušná opatření.

Evropská unie, třída B

Toto zařízení společnosti Dell je klasifikováno pro použití v obvyklém prostředí domácností (třída B).

"Prohlášení o shodě" v souladu s výše uvedenými směrnicemi a normami bylo zpracováno a je uloženo v archivu společnosti Dell™ Inc. Products Europe BV, Limerick, Irsko.

CE-krav (Europæiske Union)

Denne produkt er i overensstemmelse med 73/23/EEC (Lavspændingsdirektiv), 89/336/EEC (EMC direktiv et) og rettelsler fra den Europæiske Union.

Europæiske Union, Klasse A

ADVARSEL OM RF-FORSTYRRELSE: Dette er et Klasse A-produkt. I et hjemligt miljø kan dette produkt medføre forstyrrelse af radiofrekvens (RF), og i det tilfælde må brugeren fortage passende foranstaltninger.

Europæiske Union, Klasse B

Denne Dell-enhed er klassificeret til anvendelse i et typisk Klasse B hjemligt miljø.

En "Overensstemmelseserklæring", som er i henhold til foregående direktiver og standarder, er udført og arkiveret hos Dell Inc. Products Europe BV, Limerick, Irland.

CE-kennisgeving (Europese Unie)

Dit product voldoet aan de eisen van 73/23/EEC (laagspanningsrichtlijn), 89/336/EEC (EMC-richtlijn) en amendementen van de Europese Unie.

Europese Unie, klasse A

RF-STORINGSWAARSCHUWING: Dit is een Klasse A-product. In een woonomgeving kan dit product radiofrequentiestoring (RF-storing) veroorzaken. Indien dit zich voordoet, moet de gebruiker de passende maatregelen nemen.

Europese Unie, klasse B

Dit Dell-apparaat is geclassificeerd voor gebruik in een typische klasse B woonomgeving.

Er is in overeenstemming met de bovenstaande richtlijnen en normen een "conformiteitsverklaring" opgesteld, welke zich in het archief bevindt bij Dell™ Inc. Products Europe BV, Limerick, Ierland.

EÜ teatis (Euroopa Liit)

Käesolev toode on kooskõlas direktiividega 73/23/EMÜ (madalpinge direktiiv), 89/336/EMÜ (elektromagnetilise ühilduvuse direktiiv) ning Euroopa Liidu muudatustega.

Euroopa Liit, klass A

RAADIOSAGEDUSHÄIRE HOIATUS: Käesolev toode kuulub A-klassi. Koduses keskkonnas võib antud toode põhjustada raadiosagedushäireid, mistõttu võib selle kasutajal osutada vajalikuks võtta asjakohaseid meetmeid.

Euroopa Liit, klass B

Käesolev Dell'i seade on klassifitseeritud kasutamiseks tavapärasel klassil B vastavas olmekeskkonnas.

Vastavalt eelnevatele direktiividele ja standarditele on koostatud "vastavusdeklaratsioon", mida säilitatakse ettevõttes Dell™ Inc. Products Europe BV, Limerickis, Iirimaal.

CE-ilmoitus (Euroopan unioni)

Tämä tuote täyttää direktiivin 73/23/ETY (pienjännitedirektiivi) ja direktiivin 89/336/ETY (sähkömagneettisesta yhteensopivuudesta annettu direktiivi), sellaisina kuin ne ovat muutettuina, vaatimukset.

Euroopan unioni, Luokka A

RADIOTAAJUUSHÄIRIÖITÄ KOSKEVA VAROITUS: Tämä on Luokan A tuote. Asuinympäristössä tämä laite saattaa aiheuttaa radiotaajuushäiriöitä. Häiriöiden poiston edellyttämistä toimista vastaa laitteen käyttäjä.

Euroopan unioni, luokka B

Tämä Dell-laitte on luokiteltu käytettäväksi tyyppillisessä luokan B asuinympäristössä.

Yllä mainittujen direktiivien ja normien mukainen yhdenmukaisuusilmoitus on tehty, ja sitä säilyttää Dell™ Inc. Products Europe BV, Limerick, Irlanti.

Réglementation CE (Union européenne)

Ce produit a été déclaré conforme aux directives 73/23/EEC (Directive sur la faible tension), 89/336/EEC (Directive EMC) et aux amendements de l'Union européenne.

Union européenne, classe A

AVERTISSEMENT SUR LES PERTURBATIONS RF : Ce produit est un produit de classe A. Dans un environnement résidentiel, ce produit peut provoquer des perturbations radioélectriques, auquel cas l'utilisateur peut se voir obligé de prendre les mesures appropriées.

Union européenne, classe B

Cet appareil Dell est classé pour une utilisation dans un environnement résidentiel (classe B).

Une «Déclaration de Conformité» relative aux normes et directives précédentes a été rédigée et est enregistrée à Dell™ Inc. Products Europe BV, Limerick, Irlande.

CE-Hinweis (Europäische Union)

Es ist befunden worden, dass dieses Produkt in Übereinstimmung mit 73/23/EEC (Niederspannungs-Richtlinie), 89/336/EEC (EMC-Richtlinie) und Ergänzungen der Europäischen Union steht.

Europäische Union, Klasse A

HF-INTERFERENZWARNUNG: Dieses Produkt ist ein Produkt der Klasse A. In einer häuslichen Umgebung kann dieses Produkt Hochfrequenzstörungen verursachen. In diesem Fall muss der Benutzer die entsprechenden Maßnahmen treffen.

Europäische Union, Klasse B

Dieses Gerät von Dell ist für die Verwendung in einer typisch häuslichen Umgebung der Klasse B vorgesehen.

Eine „Konformitätserklärung“ in Übereinstimmung mit den oben angeführten Normen ist abgegeben worden und kann bei Dell Inc. Products Europe BV, Limerick, Irland, eingesehen werden.

Σήμα CE (Ευρωπαϊκή Ένωση)

Το προϊόν αυτό συμμορφώνεται με τις οδηγίες 73/23/ΕΟΚ (Οδηγία περί χαμηλής τάσης), 89/336/ΕΟΚ (Οδηγία περί ηλεκτρομαγνητικής συμβατότητας), και τροποποιήσεις τους από την Ευρωπαϊκή Ένωση.

Ευρωπαϊκή Ένωση, Κατηγορία Α

ΠΡΟΕΙΔΟΠΟΙΗΣΗ ΠΑΡΕΜΒΟΛΗΣ RF: Αυτό είναι ένα προϊόν κατηγορίας Α. Σε οικιακό περιβάλλον, αυτό το προϊόν μπορεί να προκαλέσει παρεμβολές ραδιοσυχνότητας (RF), στην οποία περίπτωση μπορεί να απαιτηθεί η λήψη κατάλληλων μέτρων από το χρήστη.

Ευρωπαϊκή Ένωση, Κατηγορία Β

Αυτή η συσκευή Dell είναι κατάλληλη για χρήση σε σύνθετες οικιακό περιβάλλον κατηγορίας Β.

Μια "Δήλωση συμμόρφωσης" βάσει των προηγούμενων Κοινοτικών Οδηγιών και Προτύπων έχει συνταχθεί και είναι αρχειοθετημένη στην Dell™ Inc. Products Europe BV, Limerick, Ιρλανδία.

CE jelzés (Európai Unió)

A termék megfelel az Európai Unió 73/23/EEC számú, kifeszültségű berendezésekre vonatkozó irányelvének, valamint a 89/336/EEC számú EMC irányelvnek és azok módosításainak.

Európai Unió, „A” osztály

RF INTERFERENCIA FIGYELMEZTETÉS: „A” osztályba sorolt termék. Lakóhelyi környezetben ez a termék rádiófrekvenciás (RF) interferenciát okozhat, ebben az esetben a felhasználónak gondoskodnia kell a szükséges ellenintézkedésekről.

Európai Unió, „B” osztály

Ez a Dell eszköz „B” osztályú besorolást kapott, tipikus lakóhelyi környezetben való használatra alkalmas.

A vonatkozó irányelvekkel és szabványokkal összhangban "Megfelelőségi nyilatkozat" készült, amely a Dell™ vállalat írországi székhelyén rendelkezésre áll (Dell Inc. Products Europe BV, Limerick, Ireland).

Avviso CE (Unione Europea)

Questo prodotto è stato determinato essere conforme alle Direttive 73/23/CEE (Direttiva sulla bassa tensione), 89/336/CEE (Direttiva CEM) ed emendamenti dell'Unione Europea.

Unione Europea, Classe A

AVVISO DI INTERFERENZA RF: Questo prodotto è classificato come Classe A. L'utilizzo di questa apparecchiatura in un'area residenziale potrebbe causare interferenze in radiofrequenza, nel qual caso potrebbe essere richiesto all'utente di intraprendere un'azione correttiva.

Unione Europea, Classe B

Il presente palmare Dell è classificato per l'uso in ambiente residenziale di Classe B.

Una "Dichiarazione di conformità" secondo gli standard e le direttive precedenti è stata emessa e registrata presso Dell™ Inc. Products Europe BV, Limerick, Irlanda.

CE atbilstības marķējums (Eiropas Savienība)

Šis produkts atbilst 73/23/EEK (Zemsprieguma Direktīva), 89/336/EEK (Elektromagnētiskās saderības Direktīva) un citiem Eiropas Savienības grozījumiem.

Eiropas Savienība, A klase

BRĪDINĀJUMS PAR RF TRAUCĒJUMIEM: Šis ir A klases produkts. Mājsaimniecības vidē produkts var radīt radio frekvenču (RF) traucējumus; šādā gadījumā lietotājam jāveic atbilstoši pasākumi.

Eiropas Savienība, B klase

Šī Dell ierīce tiek klasificēta kā izmantojama ierastos B klases mājas apstākļos.

Saskaņā ar iepriekšminētajām direktīvām un standartiem sastādīts "Atbilstības apliecinājums" un tas atrodams Dell™ Inc. Products Europe BV, Limerick, Īrijā arhīvā.

CE pranešimas (Europos Sąjunga)

Nustatyta, kad šis gaminys atitinka 73/23/EEC (žemosios įtampos direktyvą), 89/336/EEC (EMC direktyvą) ir Europos Sąjungos pataisas.

Europos Sąjunga, A klasė

ISPĖJIMAS DĖL RADIJO DAŽNIŲ TRUKDŽIŲ: Šis gaminys yra A klasės. Namų ūkio slygomis šis gaminys gali generuoti radijo dažnių trukdžius, dėl kurių vartotojas gali būti priverstas imtis atitinkamų priemonių.

Europos Sąjunga, B klasė

Šis „Dell“ įrenginys klasifikuotas kaip tinkantis naudoti tipiškoje namų ūkio (B klasės) aplinkoje.

„Atitikties deklaracija“ sukurta remiantis aukščiau išvardytomis direktyvomis bei standartais ir yra laikoma bendrovės kartotekoje adresu: Dell™ Inc. Products Europe BV, Limerick, Ireland (Airija).

Aviz CE (Unjoni Ewropea)

Ġie stabbilit li dan il-prodott hu konformi ma' 73/23/KEE (Direttiva tal-Vultaġġ Baxx), 89/336/KEE (Direttiva EMC), u emendi ta' l-Unjoni Ewropea.

Unjoni Ewropea, Klassi A

TIWISSIJA DWAR INTERFERENZA RF: Dan huwa prodott ta' Klassi A. F'ambjent domestiku dan il-prodott jista' jikkawza interferenza tal-frekwenza tar-radju (RF), fl-iema każ l-utent jista' jkun mehtieġ li jichu miżuri adegwati.

Unjoni Ewropea, Klassi B

Dan it-tagħmir Dell hu kklassifikat għall-użu f'ambjent domestiku tipiku ta' Klassi B.

Saret "Dikjarazzjoni ta' Konformità" b'konformità mad-direttivi u ma' l-istandards imsemmijin qabel. Din tinsab iffajljata għand Dell™ Inc. Products Europe BV, Limerick, l-Irlanda.

Aviso da CE (União Europeia)

Foi determinado que este produto está em conformidade com Directiva 73/23/EEC (referente a equipamentos de baixa tensão), Directiva 89/336/EEC (directiva europeia sobre compatibilidade eletromagnética) e alterações da União Europeia.

União Europeia, Classe A

ADVERTÊNCIA DE INTERFERÊNCIA DE RF: Este é um produto Classe A. Num ambiente doméstico este produto pode provocar interferência de rádio frequência (RF), podendo o utilizador ser solicitado a tomar as medidas adequadas.

União Europeia, Classe B

Este dispositivo Dell está classificado para utilização num ambiente doméstico típico Classe B.

Uma "Declaração de Conformidade" de acordo com as directivas e padrões precedentes foi elaborada e encontra-se arquivada na Dell™ Inc. Products Europe BV, Limerick, Irlanda.

Świadectwo CE (Unia Europejska)

Niniejszy produkt został uznany za zgodny z 73/23/EWG (Dyrektywą niskonapięciową), 89/336/EWG (Dyrektywą w sprawie kompatybilności elektromagnetycznej) oraz zmianami Unii Europejskiej.

Unia Europejska, klasa A

OSTRZEŻENIE O ZAKŁÓCENIACH W PAŚMIE CZĘSTOTLIWOŚCI RADIOWYCH: Urządzenie to jest urządzeniem klasy A. W środowisku domowym produkt ten może powodować zakłócenia w odbiorze fal radiowych. W takim przypadku może być konieczne podjęcie odpowiednich działań.

Unia Europejska, klasa B

Niniejsze urządzenie firmy Dell zostało zakwalifikowane do klasy B, do użytku w typowych środowiskach domowych.

"Świadectwo zgodności" zostało sporządzone zgodnie z powyższymi dyrektywami oraz normami i znajduje się w aktach firmy Dell™ Inc. Products Europe BV, Limerick, Irlandia.

CE Poznámka (Európska únia)

Tento výrobok vyhovuje požiadavkám smernice 73/23/EHS (smernica o nízkom napätí), 89/336/EHS (smernica o elektromagnetickej kompatibilite) a neskorším zmenám a doplnkom Európskej únie.

Európska únia, Trieda A

RF INTERFERENČNÉ UPOZORNENIA : Toto je zariadenie triedy A. Toto zariadenie môže v domácom prostredí spôsobiť rádiovú interferenciu, ktorú budete musieť odstrániť na vlastné náklady.

Európska únia, Trieda B

Toto zariadenie Dell triedy B je určené pre domáce prostredie.

„Vyhlásenie o zhode“ v súlade s doterajšími smernicami a normami je k dispozícii v spoločnosti Dell™ Inc. Products Europe BV, Limerick, Írsko.

Aviso CE (Unión Europea)

Este producto se ha fabricado de conformidad con la Directiva para bajo voltaje 73/23/EEC (Low Voltage Directive), la Directiva para compatibilidad electromagnética (EMC)89/336/EEC (EMC Directive), y las enmiendas de la Unión Europea.

Unión Europea, Clase A

ADVERTENCIA DE INTERFERENCIA RF: éste es un producto de Clase A. En un entorno doméstico este producto puede causar interferencia de radio frecuencia (RF), en cuyo caso el usuario debe tomar las medidas oportunas.

Unión Europea, Clase B

Este dispositivo Dell está clasificado para ser utilizado en un entorno doméstico convencional de Clase B.

Se ha realizado una "Declaración de conformidad" de acuerdo con las directivas y estándares anteriores y está archivada en Dell™ Inc. Products Europe BV, Limerick, Irlanda.

Obvestilo CE (Evropska unija)

Ta izdelek je skladen z direktivama 73/23/EGS (direktiva o nizki napetosti) in 89/336/EGS (direktiva o elektromagnetni združljivosti) ter dopolnili Evropske unije.

Evropska unija, razred A

OPOZORILO O RADIOFREKVENČNIH MOTNJAH: To je izdelek razreda A. Ta izdelek lahko v bivalnem okolju povzroča radiofrekvenčne motnje, tako da bo uporabnik moral ustrezno ukrepati.

Evropska unija, razred B

Ta Dellova naprava je razvrščena za uporabo v značilnem bivalnem okolju razreda B.

Podana je bila »Izjava o skladnosti«, skladna s prejšnjimi direktivami in standardi in je na voljo pri Dell Inc. Products Europe BV, Limerick, Irska.

CE-föreskrifter (Europeiska unionen)

Denna produkt överensstämmer med 73/23/EEC (lågspänningsdirektivet), 89/336/EEC (EMC-direktivet) och ändringar av dessa av den europeiska unionen.

Europeiska unionen, klass A

VARNING FÖR RF-STÖRNINGAR: Detta är en klass A-produkt. I bostadsmiljö kan produkten orsaka radiofrekvensstörningar. I förekommande fall måste användaren vidta lämpliga åtgärder.

Europeiska unionen, klass B

Den här Dell-enheten är klassificerad för användning i vanlig klass B-bostadsmiljö.

En "Försäkran om överensstämmelse" i enlighet med de föregående direktiven och standarderna har framställts och finns registrerad hos Dell™ Inc. Products Europe BV, Limerick, Irland.

CE Bildirimi (Avrupa Birliđi)

Bu ürünün, Avrupa Birliđi'nin deđişiklikleriyle birlikte 73/23/EEC (Düşük Voltaj Direktifi) ve 89/336/EEC (EMC Direktifi) sayılı direktiflerine uyumlu olduđu saptanmıştır.

Avrupa Birliđi, A Sınıfı

RF GİRİŞİMİ UYARISI: Bu A Sınıfı bir üründür. Evlerde kullanıldığında bu ürün radyo frekansı (RF) girişimine yol açabilir. Bu durumda kullanıcının yeterli önlemi alması gerekir.

Avrupa Birliđi, B Sınıfı

Bu Dell cihazı tipik B Sınıfı ev alanları kullanımı için sınıflandırılmıştır.

Yukarıdaki direktiflere ve standartlara göre bir "Uygunluk Bildirimi" yayınlanmış ve Dell™ Inc. Products Europe BV, Limerick, İrlanda adresinde dosyalanmıştır.

Notificare CE

S-a stabilit că acest produs respectă cerințele directivei 73/23/EEC privind joasa tensiune, ale directivei 89/336/EEC privind CEM și amendamentele Uniunii Europene.

Uniunea Europeană, Clasa A

AVERTISMENT PRIVIND INTERFERENȚELE FRECVENȚEI RADIO: Acesta este un produs din clasa A. În mediul casnic, acest produs poate cauza interferență radio, caz în care utilizatorul trebuie să ia măsurile necesare.

Uniunea Europeană, Clasa B

Acest dispozitiv Dell este clasificat pentru utilizare într-un mediu casnic obișnuit de clasă B.

Conform directivelor și standardelor precedente, a fost emisă o Declarație de Conformitate care se află depusă la Dell Inc. Products Europe BV, Limerick, Irlanda.

CE означение

Този продукт отговаря на 73/23/EEC (Нисковолтова директива), 89/336/EEC (Директива за електромагнитна съвместимост) и измененията на Европейския съюз.

Европейски съюз, Клас А

ПРЕДУПРЕЖДЕНИЕ ЗА РАДИОЧЕСТОТНИ (RF) СМУЩЕНИЯ: Това е продукт от Клас А. В жилищна среда този продукт може да създаде радиочестотни смущения, в който случай потребителят ще трябва да вземе съответните мерки.

Европейски съюз, Клас В

Това устройство на Dell е класифицирано за използване в типичната за Клас В жилищна среда.

Изготвена е „Декларация за съответствие“ според горепосочените директиви и стандарти, която се съхранява в Dell Inc. Products Europe BV, Лимерик, Ирландия.

MIC Notice (Republic of Korea Only)

MIC 규정사항 (한국에 한함)

MIC, A 등급

본 장비는 상업적인 목적을 위한 EMC 등록을 획득했습니다. 잘못 판매되거나 구매한 경우에는 가정용으로 승인된 장비로 교환하십시오.



Certification No. : (A)

다음은 MIC 규정 준수에 따라 본 설명서에서 언급하고 있는 등급 A 장치에 관한 것입니다.

- 인증 수취인 : Dell Inc.
Worldwide Regulatory Compliance & Environmental Affairs
One Dell Way
Round Rock, TX 78682 USA
512-338-4400
- 장비 또는 모델명 : 이 정보에 관한 제품 레이블을 참조하십시오.
- 인증 번호 : 이 정보에 관한 제품 레이블을 참조하십시오.
- 제조국가 : 이 정보에 관한 제품 레이블을 참조하십시오.

상기 제품의 제조 날짜는 각 제품에 기재되어 있습니다.

MIC 규정사항 (한국에 한함)

MIC, B 등급

본 장비는 가정용으로 EMC 등록을 획득했으므로, 주거 지역을 포함한 모든 지역에서 사용할 수 있습니다.



Certification No. : (B)

다음은 MIC 규정 준수에 따라 본 설명서에서 언급하고 있는 등급 B 장치에 관한 것입니다.

- 인증 수취인 : Dell Inc.
Worldwide Regulatory Compliance & Environmental Affairs
One Dell Way
Round Rock, TX 78682 USA
512-338-4400
- 장비 또는 모델명 : 이 정보에 관한 제품 레이블을 참조하십시오 .
- 인증 번호 : 이 정보에 관한 제품 레이블을 참조하십시오 .
- 제조국가 : 이 정보에 관한 제품 레이블을 참조하십시오 .

상기 제품의 제조 날짜는 각 제품에 기재되어 있습니다 .

VCCI (Japan Only)

VCCI (日本のみ)

VCCI、クラス A

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI (日本のみ)

VCCI, クラス B



この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス B 情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。

CNCA (China Only)

CNCA, Class A

声明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

información NOM (únicamente para México)

La información siguiente se proporciona en el dispositivo o dispositivos descritos en este documento, en cumplimiento con los requisitos de la Norma oficial mexicana (NOM):

Exportador: Dell Inc.
One Dell Way
Round Rock, TX 78682

Importador: Dell México, S.A. de C.V.
Paseo de la Reforma 2620 - 11° Piso
Col. Lomas Altas
11950 México, D.F.

Enviar a: Dell México, S.A. de C.V.
al Cuidado de Kuehne & Nagel de México S. de R.L.
Avenida Soles No. 55
Col. Peñon de los Baños
15520 México, D.F.

Corporate Contact Details (Taiwan Only)

Pursuant to Article 11 of the Commodity Inspection Act, Dell provides the following corporate contact details for the certified entity in Taiwan for the computer products addressed by this document:

Dell B.V. Taiwan Branch
20/F, No. 218, Sec. 2, Tung Hwa S. Road,
Taipei, Taiwan

中国大陆 RoHS

根据中国大陆《电子信息产品污染控制管理办法》(也称为中国大陆 RoHS), 以下部分列出了 Dell 产品中可能包含的有毒和 / 或有害物质的名称和含量。中国大陆 RoHS 指令包含在中国信息产业部 MCV 标准: “电子信息产品中有毒有害物质的限量要求”中。

Dell 企业产品（服务器、存储设备及网络设备）

部件名称	有毒或有害物质及元素					
	铅 (Pb)	汞 (Hg)	镉 (Cd)	六价铬 (Cr VI)	多溴联苯 (PBB)	多溴联苯醚 (PBDE)
机箱 / 挡板	X	O	X	O	O	O
印刷电路部件 - PCA*	X	O	X	O	O	O
电缆 / 连接器	X	O	X	O	O	O
硬盘驱动器	X	O	X	O	O	O
光盘驱动器 (CD、DVD 等)	X	O	O	O	O	O
磁带 / 介质	X	O	O	O	O	O
软磁盘驱动器	X	O	O	O	O	O
其它 RMSD/ 介质	X	O	O	O	O	O
电源设备 / 电源适配器	X	O	X	O	O	O
电源线	X	O	X	O	O	O
机械部件 - 风扇	X	O	O	O	O	O
机械部件 - 散热器	X	O	O	O	O	O
机械部件 - 电机	X	O	O	O	O	O
机械部件 - 其它	X	O	X	O	O	O
电池	X	O	O	O	O	O
定点设备 (鼠标等)	X	O	O	O	O	O
键盘	X	O	O	O	O	O
快擦写存储器	X	O	O	O	O	O
不间断电源设备	X	O	X	O	O	O
完整机架 / 导轨产品	X	O	X	O	O	O
软件 (CD 等)	O	O	O	O	O	O

* 印刷电路部件包括所有印刷电路板 (PCB) 及其各自的离散组件、IC 及连接器。

“O” 表明该部件所含有害和有毒物质含量低于 MCV 标准定义的阈值。

“X” 表明该部件所含有害和有毒物质含量高于 MCV 标准定义的阈值。对于所有显示 X 的情况，Dell 按照 EU RoHS 采用了容许的豁免指标。

在中国大陆销售的相应电子信息产品（EIP）都必须遵照中国大陆《电子信息产品污染控制标识要求》标准贴上环保使用期限（EPUP）标签。Dell 产品所采用的 EPUP 标签（请参阅实例，徽标内部的编号适用于指定产品）基于中国大陆的《电子信息产品环保使用期限通则》标准。



Glossary

A

Adapter

An adapter enables the computer system to access peripheral devices by converting the protocol of one bus or interface to another. An adapter may also provide specialized function. For example, a RAID controller is a type of adapter that provides RAID functions. Adapters may reside on the system board or be an add-in card. Other examples of adapters include network and SCSI adapters.

Adaptive Read Ahead

Adaptive read ahead is a read policy that specifies that the controller begins using read-ahead caching if the two most recent disk accesses occurred in sequential sectors. If all read requests are random, the algorithm reverts to Non read ahead; however, all requests are still evaluated for possible sequential operation.

B

Background Initialization

Background initialization is the automatic check for media errors on physical disks. It ensures that striped data segments are the same on all physical disks in a virtual disk. The difference between a background initialization and a consistency check is that a background initialization is automatic for new virtual disks. The operation starts within five minutes after you create the disk.

Baseport

Base register of the memory address range provided by the host.

Battery Backup Unit (BBU)

The battery backup unit protects the integrity of the cached data on the controller by providing backup power if there is a complete AC power failure or a brief power outage.

BIOS

Acronym for basic input/output system. Your computer's BIOS contains programs stored on a flash memory chip. The BIOS controls the following: communications between the microprocessor and peripheral devices, such as the keyboard and the video adapter, and miscellaneous functions, such as system messages.

BIOS Configuration Utility

The BIOS Configuration Utility configures and maintains RAID disk groups and virtual disks, and manages the RAID system. Because the utility resides in the controller BIOS, its operation is independent of the operating systems on your system. The BIOS Configuration Utility, also known as Ctrl-R, is built on elements called controls. Each control performs a function. The functions include procedures you can use to configure physical disks and virtual disks.

C

Cache

Fast memory that holds recently accessed data. Using cache speeds subsequent access to the same data. It is most often applied to processor-memory access but also can be used to store a copy of data accessible over a network. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory

software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.

Caching

The process of utilizing a high speed memory buffer, referred to as a “cache,” in order to speed up the overall read or write performance. This cache can be accessed at a higher speed than a disk subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent disk sectors. To improve write performance, the cache may temporarily store data in accordance with its write back policies. See the definition of Write-Back for more information.

Consistency Check

An operation to verify that all stripes in a virtual disk with a redundant RAID level are consistent and automatically fix any errors. For RAID 5, 10, and 50 arrays, consistency check verifies correct parity data for each stripe. For RAID 1 arrays, this operation verifies correct mirror data for each stripe.

Controller

A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a physical disk or the keyboard. In Storage Management, the hardware or logic that interacts with storage devices to write and retrieve data and perform storage management. RAID controllers perform RAID functions such as striping and mirroring to provide data protection.

D

DDR SDRAM

Acronym for Double Data Rate Synchronous Dynamic Random Access Memory. This is a type of SDRAM that provides data throughput at double the rate of conventional SDRAM. It uses a bursting technique to predict the address of the next memory location to be accessed and allows two data transfers on each clock cycle.

Disk

A non-volatile, randomly addressable, rewriteable mass storage device, including both rotating magnetic and optical storage devices and solid-state storage devices, or non-volatile electronic storage elements.

Disk Array

A set of physical disks attached to a controller or adapter.

Disk Group

A logical grouping of disks attached to a RAID controller on which one or more virtual disks can be created, such that all virtual disks in the disk group use all of the physical disks in the disk group.

Disk Migration

Moving a virtual disk or a hot spare from one controller to another by detaching the physical disks and re-attaching them to the new controller.

Disk Roaming

Moving disks from one slot to another on a controller.

Disk Subsystem

A collection of disks and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the disks can attach directly to a system I/O bus controller.

Distributed Parity

Parity involves an extra bit added to a byte or word to reveal errors in storage (in RAM or disk) or transmission. Parity is used to generate a set of redundancy data from two or more parent data sets. The redundancy data can be used to rebuild one of the parent data sets. In distributed parity, the parity data are distributed among all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity of the applicable data on the remaining physical disks.

DKMS

Acronym for Dynamic Kernel Module Support. Designed by Dell™, DKMS creates a framework in which kernel-dependent module source can reside so that it is easy to rebuild modules as you upgrade kernels. DKMS is used in the upgrade process for drivers for Red Hat® Linux and SUSE® Linux Enterprise Server.

DUD

Acronym for driver update diskette. A DUD is an image of a diskette stored as a regular file. To use it, you have to create a real diskette from this file. The steps used to create the diskette depend on how the image is supplied.

E

ECC Errors

Acronym for error correcting code. ECC errors occur in the memory and can corrupt cached data so that it must be discarded. Single-bit ECC errors can be handled by the firmware and do not disrupt normal operation. A notification will be sent if the number of single-bit errors exceeds a threshold value. ECC double-bit errors are more serious, as they result in corrupted data and data loss.

Enclosure

A structure, such as a system, which contains physical disks that are grouped together to create virtual disks.

Enclosure Management

Intelligent monitoring of the disk subsystem by software and/or hardware. The disk subsystem can be part of the host system or can reside in an external disk enclosure. Enclosure management helps you stay informed of events in the disk subsystem, such as a physical disk or power supply failure. Enclosure management increases the fault tolerance of the disk subsystem.

Exclusive-OR

A Boolean operation used to create a parity bit that can be used to restore data affected by a damaged file or failed physical disk. The management utility compares data from two physical disks and creates a parity bit that is stored on a third physical disk. This operation is used for RAID levels that use parity bits, such as RAID 5, which used distributed parity. Also known as X-OR.

F

Failed Physical Disk

A physical disk that has ceased to function, that consistently functions improperly, or that is inaccessible.

Fault Tolerance

Fault tolerance is the capability of the disk subsystem to undergo a single drive failure per disk group without compromising data integrity and processing capability. The PERC 5 controllers provide this support through redundant virtual disks in RAID levels 1, 5, 10 and 50.

Fault tolerance is often associated with system availability because it allows the system to be available during drive failures. In case a disk fails, the PERC 5 controllers support hot spare disks and the auto-rebuild feature.

Firmware

Software stored in read-only memory (ROM) or Programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a system that loads the full operating system from disk or from a network and then passes control to the operating system.

Foreign Configuration

A RAID configuration that already exists on a replacement physical disk that you install in a system. You can import the existing configuration to the RAID controller or clear it so you can create a new one.

Format

The process of writing a specific value to all data fields on a physical disk, to map out unreadable or bad sectors. Because most physical disks are formatted when manufactured, formatting is usually done only if a physical disk generates many media errors.

G

GB

Acronym for gigabyte(s). A gigabyte equals 1,024 megabytes or 1,073,741,824 bytes (2^{30} bytes).

H

Host System

Any system on which the RAID controller is installed. Mainframes, workstations, and personal systems can all be considered host systems.

Hot Spare

An idle, powered on, stand-by physical disk ready for immediate use in case of disk failure. It does not contain any user data. A hot spare can be dedicated to a single redundant virtual disk or it can be part of the global hot-spare pool for all virtual disks controlled by the controller.

When a disk fails, the controllers' firmware automatically replaces and rebuilds the data from the failed physical disk to the hot spare. Data can be rebuilt only from virtual disks with redundancy (RAID levels 1, 5, 10, or 50; not RAID 0), and the hot spare must have sufficient capacity.

Hot Swap

Replacement of a failed component while the system is running and operating normally.

I

Initialization

The process of writing zeros to the data fields of a virtual disk and, in fault tolerant RAID levels, generating the corresponding parity to put the virtual disk in a Ready state. Initializing erases previous data and generates parity so that the virtual disk will pass a consistency check. Virtual disks can work without initializing, but they can fail a consistency check because the parity fields have not been generated.

Inter-IC

Inter-IC, also known as I²C, is a multi-master bus, meaning that more than one chip can be connected to the same bus. Each chip can act as a master and initiate a data transfer.

M

MB

Acronym for megabyte(s). The term *megabyte* means 1,048,576 bytes (2²⁰ bytes); however, when referring to hard drive storage, the term is often rounded to mean 1,000,000 bytes.

Mirroring

The process of providing complete redundancy using two physical disks, by maintaining an exact copy of one physical disk's data on the second physical disk. If one physical disk fails, the contents of the other physical disk can be used to maintain the integrity of the system and to rebuild the failed physical disk.

Multi-bit ECC Errors

ECC errors are errors that occur in the memory, which can corrupt cached data so that it has to be discarded. ECC double-bit errors are serious, as they result in corrupted data and data loss. In case of double-bit ECC errors, contact Dell Technical Support.

N

Non Read Ahead

Non read ahead is a cache read policy. If you select Non read ahead in the BIOS Configuration Utility, the controller does *not* read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the data will be needed soon. Non read ahead is most effective when accessing random data.

Non Redundant Virtual Disk

A non-redundant virtual disk is one which does not have redundant data on physical disks that can be used to rebuild a failed physical disk. A RAID 0 virtual disk consists of data striped across the physical disks, without disk mirroring or parity to provide redundancy. This provides for high data throughput but offers no protection in case of a physical disk failure.

Ns

Acronym for nanosecond(s), one billionth of a second.

NVRAM

Acronym for non-volatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store configuration data on the RAID controller.

O

Offline

A physical disk is offline when it is part of a virtual disk but its data is not accessible to the virtual disk.

Online

An online device is a device that is accessible.

Online Capacity Expansion

Operation to add capacity to an existing virtual disk by adding an additional physical disk while the host system is active, and without affecting data availability.

Operating Environment

An operating environment can include the host system where physical disks are attached, any I/O buses and controllers, the host operating system and any additional software required to manage the virtual disk. For host-based arrays, the operating environment includes I/O driver software for the member disks but does not include array management software, which is regarded as part of the array itself.

P

Parity

An extra bit added to a byte or word to reveal errors in storage (in RAM or disk) or transmission. Parity is used to generate a set of redundancy data from two or more parent data sets. The redundancy data can be used to rebuild one of the parent data sets. However, parity data does not fully duplicate the parent data sets. In RAID, this method is applied to entire physical disks or stripe elements across all physical disks in a virtual disk. Parity consists of dedicated parity, in which the parity of the data on two or more physical disks is stored on an additional physical disk, and distributed parity, in which the parity data are distributed among all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity of the applicable data on the remaining physical disks.

Partition

A logical structure on a contiguous segment of storage on a physical disk or virtual disk recognized by an operating system.

Patrol Read

A preventive measure that includes review of your system for possible physical disk errors that could lead to drive failure and damage data integrity. The Patrol Read operation can find and possibly resolve any potential problem with physical disks prior to host access. This can enhance overall system performance because error recovery during a normal I/O operation may not be necessary.

PHY

The interface required to transmit and receive data packets transferred across the serial bus.

Each PHY can form one side of the physical link in a connection with a PHY on a different Dell-qualified SATA device. The physical link contains four wires that form two differential signal pairs. One differential pair transmits signals, while the other differential pair receives signals. Both differential pairs operate simultaneously and allow concurrent data transmission in both the receive and the transmit directions.

Physical Disk

A non-volatile, randomly addressable device for storing data. Physical disks are rewritable and commonly referred to as disk drives.

Physical Disk States

A physical disk can be in one of the following states:

- Un-configured Good: A disk accessible to the RAID controller but not configured as a part of a virtual disk or as a hot spare.
- Hot Spare: A physical disk that is configured as a hot spare.
- Online: A physical disk can be accessed by the RAID controller and will be part of the virtual disk.
- Rebuild: A physical disk to which data is being written to restore full redundancy for a virtual disk.

- **Failed:** A physical disk that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
- **Un-configured Bad:** A physical disk on which the firmware detects an unrecoverable error; the physical disk was Un-configured Good or the physical disk could not be initialized.
- **Missing:** A physical disk that was Online, but which has been removed from its location.
- **Offline:** A physical disk that is part of a virtual disk but which has invalid data as far as the RAID configuration is concerned.
- **None:** A physical disk with the unsupported flag set. An Un-configured Good or Offline physical disk that has completed the prepare for removal operation.

Protocol

A set of formal rules describing how to transmit data, generally across a network or when communicating with storage subsystems. Low-level protocols define the electrical and physical standards to be observed, bit- and byte-ordering, and the transmission and error detection and correction of the bit stream. High-level protocols deal with the data formatting, including the message syntax, the terminal to system dialogue, character sets, sequencing of messages, etc.

R

RAID

Acronym for Redundant Array of Independent Disks (originally Redundant Array of Inexpensive Disks). It is an array of multiple independent physical disks managed together to yield higher reliability and/or performance exceeding that of a single physical disk. The virtual disk appears to the operating system as a single storage unit. I/O is expedited because several disks can be accessed simultaneously. Redundant RAID levels (RAID levels 1, 5, 10, and 50), provide data protection.

RAID Level Migration

RAID level migration (RLM) changes the array from one RAID level to another. It is used to move between optimal RAID levels. You can perform a RLM while the system continues to run, without having to reboot. This avoids downtime and keeps data available to users.

RAID Levels

A set of techniques applied to disk groups to deliver higher data availability, and/or performance characteristics to host environments. Each virtual disk must have a RAID level assigned to it.

RAID Management Utility

A RAID management utility is used to configure physical disks into disk groups and virtual disks. The BIOS Configuration Utility is also known as Ctrl-R. Use the BIOS Configuration Utility if no operating system has been installed yet on the controller. The BIOS Configuration Utility is built on elements called controls. Each control performs a function. The functions include procedures you can use to configure physical disks and virtual disks.

The Dell OpenManage™ Storage Management configures the disks after you have installed the operating system. Storage Management enables you to perform controller and enclosure functions for all supported RAID and non-RAID controllers and enclosures from a single graphical or command-line interface without requiring the use of the controller BIOS utilities.

SAS RAID Storage Manager configures, monitors, and maintains the PERC 5 controllers, battery backup units, and other devices running on a server.

Read-Ahead

A memory caching capability in some controllers that allows them to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the additional data will be needed soon. Read-ahead supplies sequential data faster, but is not as effective when accessing random data.

Rebuild

The regeneration of all data to a replacement disk in a redundant virtual disk (RAID level 1, 5, 10, or 50) after a physical disk failure. A disk rebuild normally occurs without interrupting normal operations on the affected virtual disk, though some degradation of performance of the disk subsystem can occur.

Rebuild Rate

The percentage of central processing unit (CPU) resources devoted to rebuilding.

Reconstruct

The act of remaking a virtual disk after changing RAID levels or adding a physical disk to an existing virtual disk.

Redundancy

The provision of multiple interchangeable components to perform a single function to cope with failures and errors. Common forms of hardware redundancy are disk mirroring, implementations of parity disks, or distributed parity.

Redundant Virtual Disk

A redundant virtual disk is one which has redundant data on physical disks in the disk group that can be used to rebuild a failed physical disk. A virtual disk can use disk striping across the physical disks, disk mirroring or parity to provide redundancy. This offers protection in case of a physical disk failure.

Replacement Disk

A physical disk replacing a failed member disk in a virtual disk.

Replacement Unit

A component or collection of components in a system or subsystem that is always replaced as a unit when any part of the collection fails. Typical replacement units in a disk subsystem include disks, controller logic boards, power supplies and cables.

RPM

Acronym for Red Hat Package Manager. RPM is a software manager used to install, remove, query, and verify the software on your system. RPMs are used in the driver update procedures for Red Hat Enterprise Linux and SUSE LINUX Enterprise Server (SLES).

S**SAS**

Acronym for Serial Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the proven Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI. PERC 5 controllers leverage a common electrical and physical connection interface that is compatible with Serial ATA technology.

SATA

Acronym for Serial Advanced Technology Attachment. A physical storage interface standard, is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.

SCSIport

The SCSIport driver is a feature of the Microsoft® Windows® storage architecture, delivering SCSI commands to the storage targets. The SCSIport driver works well with storage using parallel SCSI.

Single Bit ECC Errors

ECC stands for error correcting code. ECC errors are errors that occur in the memory, which can corrupt cached data so that it has to be discarded. Single-bit ECC errors can be handled by the firmware and do not disrupt normal operation. A notification will be sent if the number of single-bit errors exceeds a threshold value.

SMART

Acronym for Self-Monitoring Analysis and Reporting Technology. The self-monitoring analysis and reporting technology (SMART) feature monitors the internal performance of all motors, heads, and drive electronics to detect predictable drive failures. This feature helps monitor drive performance and reliability, and protects the data on the drive. When problems are detected on a drive, you can replace or repair the drive without losing any data.

SMART-compliant disks have attributes for which data (values) can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical failures and some electrical failures display some degradation in performance before failure

SMP

Acronym for Serial Management Protocol. SMP communicates topology management information directly with an attached SAS expander device. Each PHY on the controller can function as an SMP initiator.

Span

A RAID technique used to combine storage space from groups of physical disks into a RAID 10 or 50 virtual disk.

Spare

A physical disk available to replace another physical disk in case that physical disk fails.

SSP

Acronym for Serial SCSI Protocol. SSP enables communication with other SAS devices. Each PHY on the SAS controller can function as an SSP initiator or SSP target.

Storport

The Storport driver has been designed to replace SCSIport and work with Windows 2003 and beyond. In addition, it offers better performance for RAID controllers, providing higher I/O throughput rates, improved manageability, and an upgraded miniport interface. For example, while SCSIport allows a maximum of 254 commands per controller, Storport allows 254 command per logical unit number (LUN).

STP

Acronym for Serial Tunneling Protocol, STP, enables communication with a Dell-qualified SATA device through an attached expander. Each PHY on the SAS controller can function as an STP initiator.

Stripe Element

A stripe element is the portion of a stripe that resides on a single physical disk. See also *striping*.

Stripe Element Size

The total disk space consumed by a stripe not including a parity disk. For example, consider a stripe that contains 64 KB of disk space and has 16 KB of data residing on each disk in the stripe. In this case, the stripe element size is 16 KB and the stripe size is 64 KB.

Striping

Disk striping writes data across all physical disks in a virtual disk. Each stripe consists of consecutive virtual disk data addresses that are mapped in fixed-size units to each physical disk in the virtual disk using a sequential pattern. For example, if the virtual disk includes five physical disks, the stripe writes data to physical disks one through five without repeating any of the physical disks. The amount of space consumed by a stripe is the same on each physical disk. The portion of a stripe that resides on a physical disk is a *stripe element*. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.

T

TBBU

Acronym for Transportable Battery Backup Unit. The TBBU protects the integrity of the cached data on the controller by providing backup power if there is a complete AC power failure or a brief power outage. A transportable battery backup unit can be used to move a controller's cached data that has not been written to the disk to a replacement controller. After you install the transportable battery backup unit on the new controller, it flushes the unwritten data preserved in the cache to the disk through the new controller.

V

Virtual Disk

A virtual disk refers to storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is seen by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data in case of a disk failure.

W

Write-Back

In write-back caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a disk write transaction. Data is written to the disk subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, elapsed time from the last cache flush, and others.

Write-Through

In write-through caching mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data and has completed the write transaction to the disk.

X

XP

XP is a Microsoft Windows operating system. Released in 2001, it is built on the Windows 2000 kernel, making it more stable and reliable than previous versions of Windows. It includes an improved user interface and more mobility features, such as plug and play features used to connect to wireless networks.

Index

B

- background
 - initialization, 17, 99
 - stopping, 67
- baseport, 99
- battery
 - installing transportable battery
 - backup, 27
 - management, 19
 - removing from PERC 5/E Adapter, 31
- BIOS, 57-58, 99
- BIOS Configuration
 - Utility, 50, 57-58, 99
 - controller management, 57
 - Foreign View menu, 58
 - menu navigation controls, 50
 - menu options, 53
 - menus, 53
 - physical disk management, 56
 - rebuild menu, 57
 - starting, 50
 - virtual disk management, 54

C

- cables
 - SAS, 72
- cache, 99
- compatibility
 - with existing RAID controllers, 16

- consistency check, 56, 66, 100
- controller, 100
 - descriptions, 11
- controller management, 57

D

- dedicated hot spares
 - managing, 64
- Disk, 18
- disk groups
 - deleting, 68
- disk migration, 18
- disk mirroring, 13
- disk roaming, 18
- disk striping, 13
- display/update parameters, 56
- distributed parity, 101
- driver diskette, 39
- driver installation, 39
- drivers
 - installation, 39
 - Microsoft operating system
 - installation, 40
 - Red Hat Enterprise Linux, 43
 - SUSE LINUX Enterprise Server, 45

E

- electrostatic discharge. *See* ESD
- ESD, 10

F

- fault tolerance, 21
 - features, 21
- firmware
 - upgrading, 68
- foreign configuration, 58
- Foreign Configuration View, 58

H

- hot swap, 102
- hot swapping, 22

I

- initialization, 103

L

- LED
 - behavior patterns, 79
 - operation, 17

M

manual rebuild, 67

O

operating system support, 12

operating systems, 12

P

parity, 14, 21, 104
distributed, 101

Patrol Read, 22
behavior, 22
behavior details, 23
blocked operations, 23
configuration, 23
modes, 23

PCI
architecture, 12

PERC
features, 15
overview, 11

PERC 5
controller descriptions, 11

physical disks
actions, 57
management, 56

post error messages, 75

R

RAID, 102-103
configuration, 49
configuration and
management, 49
configuration functions, 52
configuration information, 21
definition, 12, 105
description, 12
level migration, 105
levels, 12, 56, 105
management, 49
parity, 104
summary of levels, 12

RAID levels, 105

read policy, 59

rebuild, 57
manual, 67

Red Hat Enterprise Linux, 39
creating a driver diskette, 43
driver, 43
installing the driver, 44

regulatory notices, 81

replacement disk, 106

S

safety instructions
for preventing ESD, 10

SAS, 106
cables, 72
controller descriptions, 11
devices, 11
overview, 11

SATA, 106

SCSIport, 107

SMART technology, 17

span, 107

spare, 107

Storport, 107

stripe element size, 59

T

troubleshooting, 71
general problems, 72
physical disk issues, 72
post error messages, 75
SMART error, 74

V

virtual disks
configuring, 64
degraded, 71
deleting, 67
management, 54
menu options, 56-57
parameters, 59
setting up, 57-58

W

Windows, 39
drivers, 39
updating drivers, 41
write policy, 59