

Dell PowerEdge RAID
Controller (PERC) H700
and H800

User's Guide



Notes, Cautions, and Warnings



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



CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.



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

Information in this publication is subject to change without notice.
© 2010 Dell Inc. All rights reserved.

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

Trademarks used in this text: Dell™, the DELL logo, PowerEdge™, PowerVault™, CacheCade™, and OpenManage™ are trademarks of Dell Inc. Intel® is a registered trademark of Intel Corporation in the U.S. and other countries. Microsoft®, Windows®, Windows Server®, MS-DOS®, and Windows Vista® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat Enterprise Linux®, and Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE™ are registered trademarks of Novell Inc. in the United States and other countries. VMware® is a registered trademark of VMWare, Inc. in the United States or other countries.

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

Model UCP-70, UCPM-70, and UCP-71

March 2011

Rev. A02

Contents

1	Safety Instructions	11
	General Safety Instructions	11
	When Working Inside Your System	12
	Protecting Against Electrostatic Discharge	12
	Battery Disposal	13
2	Overview	15
	PERC H700 and H800 Card Descriptions	15
	PCI Architecture	16
	Operating System Support	16
	RAID Description	17
	Summary of RAID Levels	17
	RAID Terminology.	18
3	Storage Controller Features	21
	Physical Disk Power Management.	24
	Using the SMART Feature	24
	Initializing Virtual Disks	25
	Background Initialization of Virtual Disks	25
	Full Initialization of Virtual Disks	25
	Fast Initialization of Virtual Disks.	26

Consistency Checks	26
Disk Roaming.	26
Disk Migration	27
Compatibility With Virtual Disks Created on PERC 6 and H200 Cards	28
Migrating Virtual Disks From PERC 6 or H200 to PERC H700 and H800	28
Virtual Disk Write Cache Policies.	29
Write-Back and Write-Through.	29
Conditions Under Which Write-Back is Employed	29
Conditions Under Which Write-Through is Employed.	30
Conditions Under Which Forced Write-Back With No Battery is Employed	30
Virtual Disk Read Cache Policies	30
Reconfiguring Virtual Disks	31
Fault Tolerance Features.	33
Using Persistent Hot Spare Slots	34
Physical Disk Hot Swapping	34
Failed Physical Disk Detection	34
Redundant Path With Load Balancing Support	35
Using Replace Member and Reversible Hot Spares.	35
Controller Cache Preservation	36
Cache Preservation With Battery.	37
Cache Preservation With Non-Volatile Cache (NVC)	37
Cache Data Recovery.	37
Battery Learn Cycle.	38

	Patrol Read	38
4	Installing and Configuring Hardware	41
	Installing the PERC H700 and H800 Adapters	41
	Removing the PERC H700 and H800 Adapters	46
	Removing and Installing the PERC H700 Modular Card in Dell Blade Systems	49
	Removing the DIMM From a PERC H700	51
	Installing the DIMM on a PERC H700	53
	Replacing the BBU on a PERC H700	56
	Removing the TBBU or TNVC From a PERC H800 Adapter	58
	Replacing the Battery and Battery Cable Onto the DIMM of a PERC H800 Adapter	60
	Installing the TBBU or TNVC on a PERC H800 Adapter	61
	Transferring a TBBU or TNVC Between PERC H800 Cards	63
	Setting up Redundant Path Support on the PERC H800 Adapter	64
	Reverting From Redundant Path Support to Single Path Support on the PERC H800 Adapter	67
5	Driver Installation	69
	Installing Windows Drivers	69
	Creating the Driver Media	69

Pre-Installation Requirements	70
Installing Driver During a Windows Server 2003 Operating System Installation.	71
Installing Driver During a Windows Server 2008, Windows Server 2008 R2 Installation.	72
Installing Windows Server 2008, Windows Server 2008 R2, Windows Server 2003 for a New RAID Controller.	72
Updating Existing Windows Server 2008, Windows Server 2008 R2, Windows Server 2003.	73
Installing Linux Driver	74
Installing Red Hat Enterprise Linux Operating System Using the DUD.	76
Installing SUSE Linux Enterprise Server Using the Driver Update Diskette.	77
Installing the RPM Package With DKMS Support	78
Installing Solaris Driver	79
Installing Solaris 10 on a PowerEdge System Booting From a PERC H700 or H800 Card.	80
Adding or Updating the Driver to an Existing System	80
6 Configuring and Managing RAID	81
Dell OpenManage Storage Management	81
BIOS Configuration Utility	82
Entering the BIOS Configuration Utility.	82
Exiting the Configuration Utility.	83
Menu Navigation Controls	83
Setting Up Virtual Disks.	85
Virtual Disk Management	88

Creating Virtual Disks	88
Initializing Virtual Disks	91
Checking Data Consistency	91
Importing or Clearing Foreign Configurations Using the VD Mgmt Menu	92
Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen	93
Managing Preserved Cache	96
Managing Dedicated Hot Spares.	97
Deleting Virtual Disks.	98
Deleting Disk Groups	98
Clearing the Configuration	99
BIOS Configuration Utility Menu Options.	99
Physical Disk Management	108
Setting LED Blinking	108
Creating Global Hot Spares.	109
Removing Global or Dedicated Hot Spares	109
Replacing an Online Physical Disk	110
Stopping Background Initialization.	111
Performing a Manual Rebuild of an Individual Physical Disk.	111
Controller Management	112
Enabling Boot Support	112
Enabling BIOS Stop on Error	113
Enabling Auto Import	113
Restoring Factory Default Settings.	114
7 CacheCade	115
CacheCade Virtual Disk Characteristics	115
Configuring and Managing CacheCade Virtual Disks	116

CacheCade Virtual Disk Management	116
Create CacheCade Virtual Disk	116
Delete CacheCade Virtual Disk	118
Reconfiguring CacheCade Virtual Disks	118
8 Security Key and RAID Management	121
Security Key Implementation	121
Configuring and Managing	
Secured Virtual Disks	121
BIOS Configuration Utility	
Security Menu Options	121
Security Key Management	122
Creating Secured Virtual Disks	125
Securing Pre-Existing Virtual Disks.	126
Importing or Clearing Secured Foreign	
Configurations and Secure Disk Migration.	126
Instant Secure Erase	128
Troubleshooting Security Key Errors	128
Secured Foreign Import Errors	128
Failure to Select or Configure Non	
Self-Encrypting Disks (non-SED)	129
Failure to Delete Security Key	129
Failure to Instant Secure Erase	
Task on Physical Disks	129
9 Troubleshooting	131
Post Error Messages	131
Degraded State of Virtual Disks	141
Memory Errors	141

Preserved Cache State	142
General Issues	142
Physical Disk Related Issues	143
Physical Disk Failures and Rebuild Issues	144
SMART Errors	146
Replace Member Errors	147
Linux Operating System Errors	148
Disk Carrier LED Indicators	151
A Regulatory Notices	153
B Corporate Contact Details (Taiwan Only)	155
Glossary	157
Index	175

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.



WARNING: There is a danger of a new battery exploding if it is incorrectly installed. Replace the battery only with the same or equivalent type recommended by the manufacturer. See "Battery Disposal" on page 13.



NOTE: For complete information on 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, see the *Safety, Environmental and Regulatory Information (SERI)*, *End User License Agreement (EULA)*, and *Warranty and Support Information (WSI)* that shipped with your system.

General Safety Instructions

- 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 the compartments must 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 in 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° Celsius (140° Fahrenheit). Do not attempt to open or service batteries; replace batteries only with batteries designated for the product.

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.



CAUTION: 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 connected devices.
- 2 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.
- 3 Ground yourself by touching an unpainted metal surface on the chassis before touching anything inside the system.
- 4 While you work, periodically touch an unpainted metal surface on the chassis to dissipate any static electricity that might harm internal components.

In addition, take note of the following 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, when you connect a cable, make sure 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 system. 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 system. To prevent ESD damage, you must discharge static electricity from your body before you interact with any of your system'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 system's I/O panel) before you interact with anything electronic. When connecting a peripheral (including handheld digital assistants) to your system, you should always ground both yourself and the peripheral before connecting it to the system. Additionally, as you work inside the system, 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.

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 possible that you may never need to replace them. However, should you need to replace them, see the instructions included in the section "Configuring and Managing RAID" on page 81.



NOTE: 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. The batteries too must be disposed of in a battery deposit site. For information about such batteries, see the documentation for the specific card or component.

Taiwan Battery Recycling Mark



廢電池請回收

Overview

The Dell PowerEdge RAID Controller (PERC) H700 and H800 family of cards:

- Comply with Serial-attached SCSI (SAS) 2.0 providing up to 6 Gb/sec throughput.
- Offer RAID control capabilities which include support for RAID levels 0, 1, 5, 6, 10, 50, and 60.
- Provide reliability, high performance, and fault-tolerant disk subsystem management.

PERC H700 and H800 Card Descriptions

Table 2-1. PERC H700 and H800 Card Descriptions

Card Name	Card Description
PERC H700 Adapter	Two internal x4 SAS ports and either a battery backup unit (BBU) or non-volatile cache (NVC).
PERC H700 Integrated	Two internal x4 SAS ports and either a BBU or NVC
PERC H700 Modular	One internal x4 SAS port and a BBU
PERC H800 Adapter	Two external x4 SAS ports and either a transportable battery backup unit (TBBU) or transportable non-volatile cache (TNVC)

NOTE: Each controller supports up to 64 virtual disks.

PCI Architecture

- The PERC H700 and H800 cards support a PCI-E 2.0 x8 host interface.
- The PERC H700 Modular cards support a PCI-E 2.0 x4 host interface.

Operating System Support

The PERC H700 and H800 cards support the following operating systems:

- Microsoft Windows Server 2003 R2
- Microsoft Windows Server 2008, including Hyper-V virtualization
- Microsoft Windows Server 2008 R2
- Red Hat Enterprise Linux version 5.5 and later (32-bit and 64-bit)
- Red Hat Enterprise Linux version 6.0 and later (64-bit)
- Sun Solaris10 (64-bit)
- SUSE Linux Enterprise Server version 10 SP3 and later (64-bit)
- SUSE Linux Enterprise Server version 11 SP1 and later (64-bit)
- VMware ESX and ESXi 4.0 Update 2
- VMware ESX and ESXi 4.1



NOTE: For the latest list of supported operating systems and driver installation instructions, see the system documentation at support.dell.com/manuals. For specific operating system service pack requirements, see the **Drivers and Downloads** section at support.dell.com.

RAID Description

RAID is a group of independent physical disks that provides high performance by increasing the number of disks used for saving and accessing data.

A RAID disk subsystem offers the following benefits:

- Improves I/O performance and data availability.
- Improves data throughput because several disks are accessed simultaneously. The physical disk group appears either as a single storage unit or multiple logical units to the host system.
- Improves data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining physical disks containing data or parity.



CAUTION: In the event of a physical disk failure, a RAID 0 virtual disk fails, resulting in data loss.

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. RAID 1 is good for small databases or other applications that require small capacity and 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 6 is an extension of RAID 5 and uses an additional parity block. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 provides protection against double disk failures, and failures while a single disk is rebuilding. If you are using only one array, deploying RAID 6 is more effective than deploying a hot spare disk.
- RAID 10 is a combination of RAID 0 and RAID 1, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy. RAID 10 can support up to eight spans, and up to 32 physical disks per span.

- RAID 50 is a combination of RAID 0 and RAID 5 where a RAID 0 array is striped across RAID 5 elements. RAID 50 requires at least six disks.
- RAID 60 is a combination of RAID 0 and RAID 6 where a RAID 0 array is striped across RAID 6 elements. RAID 60 requires at least eight disks.

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 in stripes of the following sizes: 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB. The stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a stripe element.

For example, in a four-disk system using only disk striping (used in RAID 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 2-1 shows an example of disk striping.

Figure 2-1. 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 complete data redundancy. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.


Disk mirroring provides complete redundancy, but is an expensive option because each physical disk in the system must be duplicated.



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

Figure 2-2 shows an example of disk mirroring.

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



Spanned RAID Levels

Spanning is a term used to describe the way in which RAID levels 10, 50, and 60 are constructed from multiple sets of basic, or simple RAID levels. For example, a RAID 10 has multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. Similarly, RAID 50 and RAID 60 combine multiple sets of RAID 5 or RAID 6 respectively with striping.


Parity Data

Parity data is redundant data that is generated to provide fault tolerance within certain RAID levels. In the event of a disk failure the parity data can be used by the controller to regenerate your data. Parity data is present for RAID 5, 6, 50, and 60.

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 2-3. Parity provides redundancy for one physical disk failure without duplicating the contents of entire physical disks.

RAID 6 combines dual distributed parity with disk striping. This level of parity allows for two disk failures without duplicating the contents of entire physical disks.

Figure 2-3. Example of Distributed Parity (RAID 5)






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

Figure 2-4. Example of Dual Distributed Parity (RAID 6)



 **NOTE:** Parity is distributed across all disks in the array.

Storage Controller Features

This section describes the features of the Dell PowerEdge RAID Controller (PERC) H700 and H800 cards such as the configuration options, disk array performance, RAID management utilities, and operating system software drivers.

The PERC H700 and H800 family of controllers support Dell-qualified serial-attached SCSI (SAS) hard drives, SATA hard drives, and solid-state drives (SSDs).



NOTE: Mixing SAS and SATA disks within a virtual disk is not supported. Also, mixing hard drives and SSDs within a virtual disk is not supported.



NOTE: Mixing disks of different speed (10,000 rpm or 15,000 rpm) and bandwidth (3 Gbps or 6 Gbps) while maintaining the same drive type (SAS or SATA) and technology (HDD or SSD) is supported.

Table 3-1 compares the hardware configurations for the PERC H700 and H800 cards.

Table 3-1. PERC H700 and H800 Card Comparisons

Specification	PERC H700 Adapter	PERC H700 Integrated	PERC H700 Modular	PERC H800 Adapter
RAID Levels	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0,1,5,6, and 10 ^a	0, 1, 5, 6, 10, 50, 60
Enclosures per Port	N/A	N/A	N/A	Up to 4 enclosures
Ports	2 x4 internal mini-SAS wide ports	2 x4 internal mini-SAS wide ports	1x4 integrated SAS wide port	2 x4 external mini-SAS wide ports
Processor	Dell adapter SAS RAID-on-Chip, 8-port with LSI 2108 chipset	Dell adapter SAS RAID-on-Chip, 8-port with LSI 2108 chipset	Dell adapter SAS RAID-on-Chip, 4-lanes with LSI 2108 chipset	Dell adapter SAS RAID-on-Chip, 8-port with LSI 2108 chipset

Table 3-1. PERC H700 and H800 Card Comparisons

Specification	PERC H700 Adapter	PERC H700 Integrated	PERC H700 Modular	PERC H800 Adapter
BBU (Backup Battery Unit)	Optional	Optional	Yes	Optional, Transportable
Non-Volatile Cache	Optional	Optional	No	Optional, Transportable
Cache Memory	512 MB DDR2 or 1 GB DDR2	512 MB DDR2 or 1 GB DDR2	512 MB Integrated DDR2	512 MB DDR2 or 1 GB DDR2
Cache Function	Write-Back, Write-Through, Adaptive Read Ahead, No-Read Ahead, Read Ahead	Write-Back, Write-Through, Adaptive Read Ahead, No-Read Ahead, Read Ahead	Write-Back, Write-Through, Adaptive Read Ahead, No-Read Ahead, Read Ahead	Write-Back, Write-Through, Adaptive Read Ahead, Read Ahead
Maximum Number of Spans per Disk Group	8	8	2	8
Maximum Number of Virtual Disks per Disk Group	16 virtual disks per disk group	16 virtual disks per disk group	16 virtual disks per disk group	16 virtual disks per disk group
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	Up to 64 virtual disks per controller
Support for x8 2.0 PCIe Host Interface	Yes	Yes	Yes, x4 PCIe 2.0 for host operation	Yes
Online Capacity Expansion	Yes	Yes	Yes	Yes

Table 3-1. PERC H700 and H800 Card Comparisons

Specification	PERC H700 Adapter	PERC H700 Integrated	PERC H700 Modular	PERC H800 Adapter
Dedicated and Global Hot Spares	Yes	Yes	Yes	Yes
Hot Swap Devices Supported	Yes	Yes	Yes	Yes
Enclosure Hot-Add ^b	N/A	N/A	N/A	Yes
Mixed Capacity Physical Disks Supported	Yes	Yes	Yes	Yes
Hardware XOR Engine	Yes	Yes	Yes	Yes
Reversible Hot Spares Supported	Yes	Yes	Yes	Yes
Redundant Path Support	N/A	N/A	N/A	Yes
Maximum number of controllers per server	1 ^c	1	1	2 ^d

a. The RAID configurations are only supported on select Dell modular systems.

b. Using the enclosure Hot-Add feature, you can hot plug enclosures to the PERC H800 adapter without rebooting the system.

c. Only one boot controller (PERC H700) is supported on a server, consult the system *User's Guide* to determine which one is the appropriate boot controller on you system.

d. Only up to two PERC H800 controller for additional storages are supported on a system, this may be further limited by the server specifications (number of PCIe slots). Consult the system *User's Guide* for specifications.



NOTE: The maximum array size is limited by the maximum number of disks per span (32), the maximum number of spans per disk group (8), and the size of the physical disks (array and disk group terms are equivalent).

Physical Disk Power Management

The PERC H700 and H800 cards can be configured to spin down certain hard drives after a set amount of time of inactivity to conserve power. This power-savings feature is disabled by default and can be enabled in the Dell OpenManage storage management application.

The power-savings feature can be enabled so that unconfigured disks, hot spares, or both are spun down. The amount of time to wait to spin down these disks can also be set. The minimum amount of time to wait that can be set is 30 minutes and the maximum is 1 day. The default is 30 minutes.

Disks that are spun down automatically, spin up when they are needed for use. When a system is rebooted, all disks spin up.

Using the SMART Feature

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. The SMART feature helps monitor physical disk performance and reliability. SMART-compliant physical disks have attributes for which data 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.

A SMART failure is also referred to as a predicted failure. There are numerous factors that relate to predicted 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. For information on physical disk status, see "Disk Roaming" on page 26.



NOTE: For detailed information on SCSI interface specifications, see t10.org and for detailed information on SATA interface specifications, see t13.org.

Initializing Virtual Disks

You can initialize the virtual disks as described in the following sections.

Background Initialization of Virtual Disks

Background Initialization (BGI) is an automated process that writes the parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks.



NOTE: You cannot disable BGI permanently. If you cancel BGI, it automatically restarts within five minutes. For information on stopping BGI, see "Stopping Background Initialization" on page 111.

You can control the BGI rate in the Dell OpenManage storage management application. Any change in the BGI rate does not take effect until the next BGI run.



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

Consistency Check (CC) and BGI perform similar functions in that they both correct parity errors. However, CC reports data inconsistencies through an event notification, but BGI does not. You can start CC manually, but not BGI.

Full Initialization of Virtual Disks

Performing a full initialization on a virtual disk overwrites all blocks and destroys any data that previously existed on the virtual disk. Full initialization of a virtual disk eliminates the need for the virtual disk to undergo a BGI. Full initialization can be performed after the creation of a virtual disk.

During full initialization, the host is not able to access the virtual disk. You can start a full initialization on a virtual disk by using the **Slow Initialize** option in the Dell OpenManage storage management application. For more information on using the **BIOS Configuration Utility** to perform a full initialization, see "Initializing Virtual Disks" on page 91.



NOTE: If the system reboots during a full initialization, the operation aborts and a BGI begins on the virtual disk.

Fast Initialization of Virtual Disks

A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2–3 seconds to complete and is recommended when you are recreating virtual disks. To perform a fast initialization using the **BIOS Configuration Utility**, see "Initializing Virtual Disks" on page 91.



NOTE: Fast Initialization is automatically executed when a virtual disk is created with Dell OpenManage storage management application.

Consistency Checks

Consistency Check (CC) is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual disks. It is recommended that you periodically run a consistency check on virtual disks.

You can manually start a CC using the **BIOS Configuration Utility** or the Dell OpenManage storage management application. To start a CC using the **BIOS Configuration Utility**, see "Checking Data Consistency" on page 91. You can schedule CC to run on virtual disks using a Dell OpenManage storage management application.

Disk Roaming

Disk roaming is moving the physical disks from one cable connection or backplane slot to another on the same controller. The controller automatically recognizes the relocated physical disks and logically places them in the proper virtual disks that are part of the disk group. You can perform disk roaming only when the system is turned off.



CAUTION: Do not attempt disk roaming during RAID level migration (RLM) or online capacity expansion (OCE). This causes loss of the virtual disk.

Perform the following steps to use disk roaming:

- 1 Turn off the power to the system, physical disks, enclosures, and system components. Disconnect power cords from the system.
- 2 Move the physical disks to desired positions on the backplane or the enclosure.
- 3 Perform a safety check. Make sure the physical disks are inserted properly.
- 4 Turn on the system.

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

Disk Migration

The PERC H700 and H800 cards support migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is in an offline state.



NOTE: The source controller must be offline prior to performing the disk migration.



NOTE: Disks cannot be migrated back to previous PERC RAID controllers.



NOTE: Importing secured virtual disks is supported as long as the appropriate key (LKM) is supplied/configured.

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



CAUTION: Do not attempt disk roaming during RLM or online capacity expansion (OCE). This causes loss of the virtual disk.

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 online while inserting the physical disks.

The controller flags the inserted disks as foreign disks.

- 3 Use the Dell OpenManage storage management application or the controller **BIOS Configuration Utility** to import the detected foreign configuration.
- 4 Ensure that all physical disks that are part of the virtual disk are migrated.



NOTE: For more information about compatibility, contact your Dell technical support representative.

Compatibility With Virtual Disks Created on PERC 6 and H200 Cards

Virtual disks that are created on the PERC 6 and H200 family of controllers can be migrated to the PERC H700 and H800 cards without risking data or configuration loss. Migrating virtual disks from the PERC H700 and H800 cards to PERC 6 or H200 is not supported.



NOTE: For more information about compatibility, contact your Dell technical support representative.

Migrating Virtual Disks From PERC 6 or H200 to PERC H700 and H800

To migrate virtual disks from PERC 6 or H200 to PERC H700 and H800:

- 1 Turn off the system.
- 2 Move the appropriate physical disks from the PERC 6 or H200 card to the PERC H700 and H800 card. If you are replacing your PERC 6 or H200 card with a PERC H700 or H800 card, see the *Hardware Owner's Manual* that shipped with your system or at support.dell.com/manuals.
- 3 Boot the system and import the foreign configuration that is detected. You can do this in two ways:
 - Press <F> to automatically import the foreign configuration.
 - Enter the **BIOS Configuration Utility** and navigate to the **Foreign Configuration View**.



NOTE: For more information on accessing the **BIOS Configuration Utility**, see "Entering the BIOS Configuration Utility" on page 82.



NOTE: For more information on **Foreign Configuration View**, see "Foreign Configuration View" on page 108.

- 4 Exit the **BIOS Configuration Utility** and reboot the system.
- 5 Ensure all the latest drivers for the PERC H700 or H800 card (available at support.dell.com) are installed. For more information, see "Driver Installation" on page 69.

Virtual Disk Write Cache Policies

The write cache policy of a virtual disk determines how the controller handles writes to that virtual disk. **Write-Back** and **Write-Through** are the two write cache policies and can be set on virtual disks individually.

All RAID volumes are presented as Write-Through (WT) to the operating system (Windows and Linux) independent of the actual write cache policy of the virtual disk. The PERC cards manage the data in cache independently of the operating system or any applications. Use Dell OpenManage or the **BIOS Configuration Utility** to view and manage virtual disk cache settings.

Write-Back and 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.

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 controller then writes the cached data to the storage device in the background.

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 mitigated by using a BBU on PERC H700 or H800 cards. For information on which controllers support a BBU, see Table 3-1.

Write-Back caching has a performance advantage over **Write-Through** caching.



NOTE: The default cache setting for virtual disks is **Write-Back** caching.



NOTE: Certain data patterns and configurations perform better with a **Write-Through** cache policy.

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. This low-charge state does not apply to controllers with the optional non-volatile cache (NVC) module present.

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

Write-Back mode is available when you select **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 not present.

 **CAUTION:** It is recommended that you use a power backup system when forcing **Write-Back** to ensure there is no loss of data if the system suddenly loses power.

Virtual Disk Read Cache Policies

The read policy of a virtual disk determines how the controller handles reads to that virtual disk. The read policies are:

- **Always Read Ahead** — Allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data is required soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **No Read Ahead** — Disables the **Read-Ahead** capability.
- **Adaptive Read Ahead** — Begins using **Read-Ahead** if the two most recent disk accesses occurred in sequential sectors. If the read requests are random, the controller reverts to **No Read Ahead** mode.

 **NOTE:** The default read cache setting for virtual disks is **Adaptive Read Ahead**.

Reconfiguring Virtual Disks

An online virtual disk can be reconfigured in ways that expands its capacity and/or change its RAID level. Spanned virtual disks such as RAID 10, 50, and 60 cannot be reconfigured.

Online Capacity Expansion (OCE) can be done in two ways. The first way is to expand the volume using free space already in the disk group. For example, if there is a single virtual disk in a disk group and free space is available, the virtual disk's capacity can be expanded within that free space. If a virtual disk is created and it does not use the maximum size of the disk group, free space is available and OCE is possible. Free space is also available when a disk group's physical disks are replaced by larger disks using the Replace Member feature. The second method by which OCE can be done is by adding physical disks to the disk group and using OCE to incorporate free space from the newly added disks into the virtual disk.

RAID Level Migration (RLM) refers to changing a virtual disk's RAID level. Both RLM and OCE can be done at the same time so that a virtual disk can simultaneously have its RAID level changed and its capacity increased. When a RLM/OCE operation is complete, a reboot is not required. For a list of RLM/OCE possibilities, see Table 3-2. The source RAID level column indicates the virtual disk RAID level before the RLM/OCE and the target RAID level column indicates the RAID level after the operation has completed.



NOTE: If you configure 64 virtual disks on a controller, you cannot perform a RAID level migration or capacity expansion on any virtual disk.



NOTE: The controller changes the write cache policy of all virtual disks undergoing a RLM/OCE to **Write-Through** until the RLM/OCE is complete.

Table 3-2. RAID Level Migration

Source RAID Level	Target RAID Level	Required Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
RAID 0	RAID 0	1	2 or more	Yes	Increases capacity by adding disks
RAID 0	RAID 1	1	2	No	Converts non-redundant virtual disk into a mirrored virtual disk by adding one disk.
RAID 0	RAID 5	1 or more	3 or more	Yes	At least one disk needs to be added for distributed parity data.
RAID 0	RAID 6	1 or more	4 or more	Yes	At least two disks need to be added for dual distributed parity data.
RAID 1	RAID 0	2	2 or more	Yes	Removes redundancy while increasing capacity.
RAID 1	RAID 5	2	3 or more	Yes	Maintains redundancy while doubling capacity.
RAID 1	RAID 6	2	4 or more	Yes	Two disks required to be added for distributed parity data.
RAID 5	RAID 0	3 or more	3 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space used for distributed parity data.
RAID 5	RAID 5	3	4 or more	Yes	Increases capacity by adding disks

Table 3-2. RAID Level Migration (continued)

Source RAID Level	Target RAID Level	Required Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
RAID 5	RAID 6	3 or more	4 or more	Yes	At least one disk needs to be added for dual distributed parity data.
RAID 6	RAID 0	4 or more	4 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space used for distributed parity data.
RAID 6	RAID 5	4 or more	4 or more	Yes	Removes one set of parity data and reclaims disk space used for it.
RAID 6	RAID 6	4	5 or more	Yes	Increases capacity by adding disks



NOTE: The total number of physical disks in a disk group cannot exceed 32. You cannot perform RAID level migration and/or expansion on RAID levels 10, 50, and 60.

Fault Tolerance Features

The list of features of the controller cards that provide fault tolerance to prevent data loss is as follows:

- Support for SMART
- Support for Patrol Read
- Redundant path support (for PERC H800 only)
- Physical disk failure detection
- Physical disk rebuild using hot spares
- Parity generation and checking (for RAID 5, 50, 6, and 60 only)

- Battery and optional Non-Volatile Cache backup of controller cache to protect data
- Detection of batteries with low charge after boot up

The next sections describe some methods to achieve fault tolerance.

Using Persistent Hot Spare Slots

The H700 and H800 cards can be configured so that system backplane or storage enclosure disk slots are dedicated as hot spare slots. This feature can be enabled using the Dell OpenManage storage management application.

Once enabled, any slots with hot spares configured automatically become persistent hot spare slots. If a hot spare disk fails or is removed, a replacement disk that is inserted into the same slot automatically becomes a hot spare with the same properties as the one it is replacing. If the replacement disk does not match the disk protocol and technology, it does not become a hot spare.



NOTE: Any hot spares assigned to an encrypted Virtual Disk must also be capable of encryption (SED).

Physical Disk Hot Swapping

Hot swapping is the manual replacement of a disk while the H700 and H800 cards are online and performing their normal functions.

The following requirements must be met before hot swapping a physical disk:

- The system backplane or enclosure must support hot swapping for the PERC H700 and H800 cards to support hot swapping.
- The replacement disk must be of the same protocol and disk technology. For example, only a SAS hard drive can replace a SAS hard drive; only a SATA SSD can replace a SATA SSD.
- The replacement disk must be of equal or greater capacity than the one it is replacing.

Failed Physical Disk Detection

Failed physical disks are detected and rebuilds automatically start to new disks that are inserted into the same slot. Automatic rebuilds can also happen transparently with hot spares. If you have configured hot spares, the controllers automatically try to use them to rebuild failed physical disks.

Redundant Path With Load Balancing Support

The PERC H800 adapter can detect and use redundant paths to disks contained in enclosures. This provides the ability to connect two SAS cables between a controller and an enclosure for path redundancy. The controller is able to tolerate the failure of a cable or Enclosure Management Module (EMM) by utilizing the remaining path.

When redundant paths exist, the controller automatically balances I/O load through both paths to each disk. Load balancing increases throughput to virtual disks in storage enclosures and is automatically turned on when redundant paths are detected. The ability to load balance I/O can be disabled using the Dell OpenManage storage management application. To set up your hardware to support redundant paths, see "Setting up Redundant Path Support on the PERC H800 Adapter" on page 64.



NOTE: The hardware support for redundant paths refer to path-redundancy only and not to controller-redundancy.

Using Replace Member and Revertible Hot Spares

The **Replace Member** functionality allows a previously commissioned hot spare to be reverted to a usable hot spare. When a disk failure occurs within a virtual disk, an assigned hot spare (dedicated or global) is commissioned and begins rebuilding until the virtual disk is optimal. After the failed disk is replaced (in the same slot) and the rebuild to the hot spare is complete, the controller automatically starts to copy data from the commissioned hot spare to the newly-inserted disk. After the data is copied, the new disk is a part of the virtual disk and the hot spare is reverted to being a ready hot spare. This allows hot spares to remain in specific enclosure slots. While the controller is reverting the hot spare, the virtual disk remains optimal.



NOTE: The controller automatically reverts a hot spare only if the failed disk is replaced with a new disk in the same slot. If the new disk is not placed in the same slot, a manual **Replace Member** operation can be used to revert a previously commissioned hot spare.

Automatic Replace Member with Predicted Failure

A **Replace Member** operation can occur when there is a SMART predictive failure reporting on a physical disk in a virtual disk. The automatic **Replace Member** is initiated when the first SMART error occurs on a physical disk that is part of a virtual disk. The target disk needs to be a hot spare that qualifies as a rebuild disk. The physical disk with the SMART error is marked as **failed** only after the successful completion of the **Replace Member**. This avoids putting the array in degraded status.

If an automatic **Replace Member** occurs using a source disk that was originally a hot spare (that was used in a rebuild), and a new disk added for the **Replace Member** operation as the target disk, the hot spare reverts to the hot spare state after a successful **Replace Member** operation.



NOTE: To enable the automatic **Replace Member**, use the Dell OpenManage storage management application. For more information on automatic **Replace Member**, see "Dell OpenManage Storage Management" on page 81.



NOTE: For information on manual **Replace Member**, see "Replacing an Online Physical Disk" on page 110.

Controller Cache Preservation

The controller is capable of preserving its cache in the event of a system power outage or improper system shutdown. The PERC H700 controllers are attached to a Battery Backup Unit (BBU) that provides backup power during system power loss to preserve the controller's cache data. The PERC H800 has a transportable version of the BBU attached to it called the Transportable Battery Backup Unit (TBBU) which enables the entire cache module to be transported to a new controller if necessary. If the card has the optional non-volatile cache (NVC) or transportable non-volatile cache (TNVC) module, the cache data is preserved using flash storage instead of battery power.

Cache Preservation With Battery

The lithium-ion battery included in the BBU/TBBU of the controller is an inexpensive way to protect data in cache memory. If the controller has data in cache memory during a power outage or improper system shutdown, battery power is used to preserve cache data until power is restored or the battery is depleted. Under the 1-year limited warranty, the battery provides at least 24 hours of backup power coverage in normal operating conditions during the warranty period. To prolong battery life, do not store or operate the battery in temperatures that exceed 60 degrees C.

Cache Preservation With Non-Volatile Cache (NVC)

The NVC module allows controller cache data to be stored indefinitely, an advantage over the 24 hours that battery backup provides. If the controller has data in cache memory during a power outage or improper system shutdown, a small amount of power from a battery is used to transfer cache data to non-volatile flash storage where it remains until power is restored and the system is booted.

Cache Data Recovery

The dirty cache LED that is located on the H700 and H800 cards cannot be used to determine if cache data is being preserved. If a system power loss or improper system shutdown has occurred, restore system power and boot the system. During the boot, enter the controller's BIOS Configuration Utility (<Ctrl><R>) to ensure that there is no cache data being preserved. This can be done by entering into the controller menu and selecting Managed Preserved Cache. If there are no virtual disks listed here, all preserved cache data has been written to disk successfully.

In the event of a PERC H800 card failure, the entire TBBU/TNVC module can be safely transferred to a new PERC H800 card without putting preserved cache data at risk. See the applicable sections in "Installing and Configuring Hardware" on page 41 for instructions on removing and installing the TBBU/TNVC, then follow the instructions above on recovering cache data.

Battery Learn Cycle

Learn cycle is a battery calibration operation performed by the controller periodically to determine the condition of the battery. This operation cannot be disabled.



NOTE: Virtual disks automatically switch to **Write-Through** mode when the battery charge is low because of a learn cycle.

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 H700 or H800 cards, the expected time frame for completion of a learn cycle is approximately seven hours and consists of the following parts:

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



NOTE: For additional information, see the Dell OpenManage storage management application.

During the discharge phase of a learn cycle, the PERC H700 or H800 battery charger is disabled and remains disabled until the battery is discharged. After the battery is discharged, the charger is re-enabled.

Patrol Read

The **Patrol Read** feature is designed as a preventative measure to ensure physical disk health and data integrity. **Patrol Read** scans for and resolves potential problems on configured physical disks. The Dell OpenManage storage management application can be used to start **Patrol Read** and change its behavior.

The following is an overview of **Patrol Read** behavior:


- **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 physical disks that are not part of a virtual disk or are in **Ready** state.


- **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** uses fewer resources to allow the I/O to take a higher priority.
- **Patrol Read** does not run on any disks involved in any of the following operations:
 - Rebuild
 - Replace Member
 - Full or Background Initialization
 - CC
 - RLM or OCE




NOTE: By default, **Patrol Read** automatically runs every seven days on configured SAS and SATA hard drives. **Patrol Read** is not necessary on SSD and is disabled by default.

Installing and Configuring Hardware


 **WARNING:** 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.

 **CAUTION:** Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

 **NOTE:** For complete information on 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, see the *Safety, Environmental and Regulatory Information (SERI)*, *End User License Agreement (EULA)*, and *Warranty and Support Information (WSI)* that shipped with your system.

Installing the PERC H700 and H800 Adapters

- 1 Unpack the Dell PowerEdge RAID Controller (PERC) 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. For more information on preparing your system for hardware changes, see the *Hardware Owner's Manual* shipped with your system or at: support.dell.com/manuals.
- 3 Disconnect all attached devices and remove the system cover. For more information on opening the system, see your system's *Hardware Owner's Manual*.
- 4 Select an empty PCI-E slot. Remove the blank filler bracket on the back of the system aligned with the PCI-E slot you have selected.


NOTE: The PERC H700 Integrated and H700 Modular cards have a dedicated storage slot. For details on the correct PCI-E location, see the *Hardware Owner's Manual* that shipped with your system or see the appropriate documentation available at support.dell.com/manuals.

5 Align the PERC H700 or H800 card to the PCI-E slot you have selected.

CAUTION: Never apply pressure to the adapter module while inserting it in the PCI-E slot. Applying pressure could break the adapter module.

6 Insert the controller gently, but firmly, until the controller is firmly seated in the PCI-E slot. For more information on installing the PERC H800 adapter, see Figure 4-1. For more information on installing the PERC H700 adapter, see Figure 4-2.


Figure 4-1. Installing a PERC H800 Adapter



1 PCI-e slot
3 filler bracket

2 PERC H800 adapter
4 bracket screw


Figure 4-2. Installing a PERC H700 Adapter



- | | | | |
|---|-----------------|---|-------------------|
| 1 | PCI-e slot | 2 | PERC H700 adapter |
| 3 | filler brackets | 4 | bracket screw |

- 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. For more information on closing the system, see the *Hardware Owner's Manual* shipped with your system or at support.dell.com/manuals.
- 9 For the PERC H800 adapter, connect the cable from the external enclosure to the controller. For more information, see Figure 4-3.


Figure 4-3. Connecting the Cable From the External Enclosure



- 1 connector on the controller
- 2 cable from the external enclosure
- 3 system

10 For the PERC H700 adapter, connect the cables from the backplane of the system to the controller. The primary SAS connector is labeled **SAS_A** and the secondary SAS connector is labeled **SAS_B**. For more information, see Figure 4-4.


Figure 4-4. Connecting Cables to the Controller




- | | | | |
|---|-----------|---|-------------------|
| 1 | cable | 2 | Port B |
| 3 | Port A | 4 | PERC H700 adapter |
| 5 | connector | | |


- 11** Replace the cover of the system. For more information on closing the system, see the *Hardware Owner's Manual* shipped with your system or at support.dell.com/manuals.
- 12** Reconnect the power and network cables, and turn on the system.

Removing the PERC H700 and H800 Adapters

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

 **NOTE:** Before beginning with the procedure, press <Ctrl><R> when the system is booting to verify that no cache is preserved.

- 1 Perform a controlled reboot of the system and enter the **PERC BIOS Configuration Utility** to ensure that there is no data present in cache. See the "Cache Data Recovery" on page 37 for more details. Later, perform a controller shutdown of the system as well as any attached storage controllers.
- 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 peripherals installed in the system's PCI-E slots, see the *Hardware Owner's Manual* shipped with your system or at support.dell.com/manuals.

For instructions on removing a PERC H800 adapter, go to step 3.
For instructions on removing a PERC H700 adapter, go to step 5.

- 3 Locate the PERC H800 adapter in the system and disconnect the external cables from the adapter.
- 4 Remove any retention mechanism, such as a bracket screw, that may be holding the PERC H800 in the system and gently lift the controller from the system's PCI-E slot. For more information, see Figure 4-5.


Figure 4-5. Removing the PERC H800 Adapter



- | | |
|---------------------|-------------------|
| 1 bracket screw | 2 memory module |
| 3 battery | 4 Dirty Cache LED |
| 5 PERC H800 Adapter | |

- 5 Disconnect the data cables and battery cable from the PERC H700. Remove any retention mechanism, such as a bracket screw, that might be holding the PERC H700 in the system, and gently lift the controller from the system's PCI-E slot.

Figure 4-6. Removing the PERC H700 Adapter



1 bracket screw

2 Dirty Cache LED

3 PERC H700 controller

Removing and Installing the PERC H700 Modular Card in Dell Blade Systems



NOTE: For more information on removing and installing blade system parts, see your system's *Hardware Owner's Manual* or the *User's Guide* from the Dell Support website at support.dell.com.

The storage controller card is located below the disk bays of the Dell Blade system.


To remove the storage controller card:

- 1 Remove the Dell Blade system from the Blade system chassis.
- 2 Remove the system cover of the Blade system.
- 3 Remove the system board and place it on a stable and flat surface.
- 4 Determine whether the dirty cache LED on the controller is illuminated. For location of the LED, see Figure 4-7.

If the LED is illuminated, re-insert the system board, replace the system cover, reconnect the system to power, turn on the system, and repeat step 1 through step 3. If the LED is not illuminated, continue with the next step.

- 5 Open the release lever to disconnect the storage controller card edge connector from the system board connector as illustrated in Figure 4-7.
- 6 Lift the storage controller card straight up from the system board as illustrated in Figure 4-7.

Figure 4-7. Removing and Installing the Storage Controller Card



- | | | | |
|---|-------------------------|---|-------------------------|
| 1 | battery cable connector | 2 | storage controller card |
| 3 | release lever | 4 | dirty cache LED |

To install your new storage controller card:

- 1 Unpack the new storage controller card and check for damage.
NOTE: If the card is damaged, contact Dell technical support.
- 2 Place the storage controller card onto the system board. Align the storage controller card such that the tabs on the system board tray fit through the notches on the edges of the storage controller card.
- 3 Slide the storage controller card towards the connector on the system board until the storage controller clicks in place.
- 4 Reinstall the system board. For more information on reinstalling the system board, see your system's *Hardware Owner's Manual* or the *User's Guide*.

- 5 Close the top cover of the Blade system. For more information on closing the top cover of the Modular Blade system, see your system's *Hardware Owner's Manual* or the *User's Guide*.
- 6 Reinstall the Blade system in the Blade system chassis. For more information on reinstalling the Blade system in the Blade system chassis, see your system's *Hardware Owner's Manual* or the *User's Guide*.



NOTE: For the latest list of firmware and installation instructions, see the Dell Support website at support.dell.com.

Removing the DIMM From a PERC H700



CAUTION: PERC H700 Modular cards, shipped in PowerEdge blade systems, have an integrated DIMM module which cannot be removed. Do not attempt the following procedure on a PERC H700 Modular controller card.

- 1 Perform a controlled reboot of the system and enter the PERC H700 BIOS Configuration Utility to ensure that there is no data present in cache. See the "Cache Data Recovery" on page 37 for more details. Later, shut down the system.



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

- 2 Disconnect the system from the electrical outlet and remove the system cover.




NOTE: The location of the PERC H700 controller varies from system to system. For more information on PERC H700 location, see the *Hardware Owner's Manual* that shipped with your system or see the appropriate documentation available at support.dell.com/manuals.

- 3 Remove the PERC H700 from the system. See "Removing the PERC H700 and H800 Adapters" on page 46.
- 4 Remove the DIMM by pressing down on the tabs at each edge of the DIMM connector and lift the DIMM off the controller. See Figure 4-8.



NOTE: Do not exert excessive pressure on the connector of the DIMM while removing the DIMM.

Figure 4-8. Removing the DIMM From a PERC H700



- | | | | |
|---|-------------------|---|-----------------|
| 1 | DIMM support | 2 | dirty cache LED |
| 3 | DIMM | 4 | tab |
| 5 | PERC H700 Adapter | | |

Installing the DIMM on a PERC H700


△ CAUTION: PERC H700 Modular cards, shipped in PowerEdge blade systems, have an integrated DIMM module which cannot be removed. Do not attempt the following procedure on a PERC H700 Modular controller card.

- 1 Unpack the DIMM and follow all antistatic procedures.

✎ NOTE: Do not use excessive pressure on the connector on the DIMM while installing the DIMM.

- 2 If the DIMM support is not mounted onto the DIMM, perform the following steps to replace the DIMM support before mounting on a PERC H700:
 - a With the old DIMM removed from the controller, press out on the DIMM support clips inserted through the DIMM rotating the DIMM support out of the DIMM. See Figure 4-9.

Figure 4-9. Removing the DIMM Support From the DIMM




1 DIMM

2 DIMM support

- b To mount, place the top edge of the DIMM support over the top edge of the DIMM so that the arms on the side of the DIMM support fit into their sockets on the DIMM. See Figure 4-10.

Figure 4-10. Mounting the DIMM Support onto the DIMM




1 DIMM

2 DIMM support

- 3** Align the keyed edge of the DIMM to the physical divider on the memory socket to avoid damage to the module.
- 4** Insert the DIMM in the memory socket. Apply a constant, downward pressure on both ends or the middle of the DIMM until the retention clips fall in the allotted slots on either side of the DIMM. See Figure 4-11.

Figure 4-11. Installing a DIMM on a PERC H700



- | | | | |
|---|-----------------|---|----------------|
| 1 | PERC H700 | 2 | retention clip |
| 3 | memory socket | 4 | DIMM |
| 5 | dirty cache LED | | |

Replacing the BBU on a PERC H700


- 1 Perform a controlled reboot of the system and enter the **PERC H700 BIOS Configuration Utility** to ensure that there is no data present in cache. See the "Cache Data Recovery" on page 37 for more details. Later, shut down the system.

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

- 2 Disconnect the system from the electrical outlet and remove the system cover.


🔧 NOTE: The locations of the PERC H700 and BBU vary from system to system. For more information on PERC H700 card and BBU locations, see the *Hardware Owner's Manual* shipped with your system or see the appropriate documentation available at support.dell.com/manuals.

Figure 4-12. Dirty Cache LED and Battery Cable Connector



- | | | | |
|---|-------------------------|---|-------------------|
| 1 | battery cable connector | 2 | PERC H700 Adapter |
| 3 | dirty cache LED | | |

Figure 4-13. PERC H700 (Modular) Dirty Cache LED Location



- | | | | |
|---|---------------------------|---|-----------------|
| 1 | battery cable connector | 2 | dirty cache LED |
| 3 | PERC H700 Modular Adapter | | |

3 Locate the battery cable connection near the edge of the controller, and disconnect the battery. For the location of the battery cable connector, see Figure 4-12 and Figure 4-13.


4 Remove the battery from the plastic mounting shroud in your system and disconnect the battery cable. For more information on the location of the BBU in your system and instructions on how to replace the BBU, see the *Hardware Owner's Manual* that shipped with the system.

NOTE: Connect the battery cable to the new BBU before mounting it to the plastic shroud in your system.

5 Re-connect the battery to the controller by inserting the end of the battery cable into the connector on the controller.


NOTE: Do not exert excessive pressure connecting the cable. It is a keyed connector and it fits only one way.

Removing the TBBU or TNVC From a PERC H800 Adapter

 **NOTE:** Both the TBBU and TNVC on the PERC H800 Adapter consists of the DIMM and a battery.

- 1 Perform a controlled reboot of the system and enter the **PERC BIOS Configuration Utility** to ensure that there is no data present in cache. See "Cache Data Recovery" on page 37 for more details. Later, perform a controller shutdown of the system as well as any attached storage controllers.
- 2 Disconnect the system from the electrical outlet and remove the system cover.
- 3 Remove the PERC H800 Adapter from the system. For more information, see "Removing the PERC H700 and H800 Adapters" on page 46.
- 4 Press down on the tabs at each edge of the DIMM slot and lift the TBBU or TNVC assembly off the PERC H800 Adapter. See Figure 4-14.

Figure 4-14. Removing the TBBU From a PERC H800 Adapter




- 1 battery
- 3 DIMM

- 2 dirty cache LED
- 4 PERC H800 Adapter

Replacing the Battery and Battery Cable Onto the DIMM of a PERC H800 Adapter

- 1 With the old TBBU or TNVC removed from the controller, disconnect the battery cable from both ends and press out on the battery clips inserted through the DIMM rotating the battery out of the DIMM. See Figure 4-15.


Figure 4-15. Removing the Battery and Battery Cable From a PERC H800 Adapter



1 DIMM


2 battery

- 2 Unpack the new TBBU or TNVC and follow all antistatic procedures.
- 3 Insert one end of the battery cable into the connector on the DIMM and the other end into the connector on the new battery.

 **NOTE:** Do not exert excessive pressure connecting the cable. It is a keyed connector and it fits only one way.


- 4 Place the top edge of the battery over the top edge of the DIMM so that the arms on the side of the battery fit into their sockets on the DIMM. See Figure 4-16.

Figure 4-16. Installing the Battery and Battery Cable onto the DIMM



- | | | | |
|---|-----------------------|---|--------------------------|
| 1 | connector on the DIMM | 2 | DIMM |
| 3 | battery | 4 | connector on the battery |
| 5 | battery cable | | |

Installing the TBBU or TNVC on a PERC H800 Adapter


 **NOTE:** Both the TBBU and TNVC on the PERC H800 Adapter consists of the DIMM and a battery.

- 1 Check the battery attachment to the DIMM and battery cable connection. For more information, see "Replacing the Battery and Battery Cable Onto the DIMM of a PERC H800 Adapter" on page 60.

 **CAUTION:** Do not touch the gold leads and do not bend the memory module.

- 2 Align the keyed edge of the DIMM to the physical divider on the memory socket to avoid damage to the module.
- 3 Insert the DIMM in the memory socket. Apply a constant, downward pressure on both ends or the middle of the DIMM until the retention clips fall in the allotted slots on either side of the DIMM. See Figure 4-17.

Figure 4-17. Installing the TBBU on a PERC H800 Adapter



- | | | | |
|---|---------------|---|----------------|
| 1 | PERC H800 | 2 | retention clip |
| 3 | memory socket | 4 | DIMM |

Transferring a TBBU or TNVC Between PERC H800 Cards

The TBBU or TNVC provides uninterrupted power supply for up to 48 hours to a cache memory module. If the controller fails as a result of a power failure, you can move the TBBU or TNVC to a new controller and recover the data. The controller that replaces the failed controller should not have any prior configuration.

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

- 1** Perform a controlled shutdown on the system in which the PERC H800 is installed, as well as any attached storage enclosures.
- 2** Disconnect the system from the electrical outlet and remove the system cover.
- 3** Remove the controller that has the TBBU or TNVC.
- 4** Remove the TBBU or TNVC from the controller.
- 5** Insert the TBBU or TNVC in the new controller.

For more information on installing the TBBU or TNVC, see "Installing the TBBU or TNVC on a PERC H800 Adapter" on page 61.

- 6** Insert the replacement controller in the system.
See the relevant sections on installing controllers under "Installing the PERC H700 and H800 Adapters" on page 41.
- 7** Reconnect all external storage enclosures as they were originally connected and turn on the enclosures.
- 8** Replace the system cover, re-connect the system to the electrical outlet and turn on the system.

The controller flushes the cache data to the virtual disks.

Setting up Redundant Path Support on the PERC H800 Adapter

The PERC H800 adapter can detect and use redundant paths to disks contained in enclosures. With redundant paths to the same device, if one path fails, another path can be used to communicate between the controller and the device. For more information about redundant paths, see "Redundant Path With Load Balancing Support" on page 35.

To set up a configuration with redundant paths, both ports on a controller must be cabled to the **In** ports of a single enclosure.

To add multiple enclosures, both **Out** ports of the first enclosure must be cabled to the **In** ports of the next enclosure.

If the connection between an **Out** port on the controller and an **In** port on an enclosure fails, an alternate path exists through the second **Out** port on the controller and the second **In** port on the enclosure. For more information, see Figure 4-18 to Figure 4-21.



NOTE: The PERC H800 adapter supports redundant paths when used with Dell PowerVault MD1200 and Dell PowerVault MD1220 disk storage enclosures.

Perform the following steps to configure the hardware to utilize redundant paths on the PERC H800 adapter:

- 1 Set up an enclosure on the PERC H800 adapter.
- 2 Connect two SAS cables from the **Out** ports on your PERC H800 adapter to the **In** ports of the external enclosure. For more information, see Figure 4-18.



NOTE: For information on Unified Mode, see the enclosure documentation that was shipped with the enclosure.

- 3 To add multiple enclosures, cable both **Out** ports of the first enclosure to both **In** ports of the next enclosure.

After you set up the hardware, the controller detects the redundant paths and automatically utilizes them to balance the I/O load.

Figure 4-18 displays redundant path storage configuration with one enclosure.

Figure 4-18. Redundant Path Support Configuration With One Enclosure




Figure 4-19 displays redundant path storage configuration with two enclosures.

Figure 4-19. Redundant Path Support Configuration With Two Enclosures




Figure 4-20 displays redundant path storage configuration with three enclosures.

Figure 4-20. Redundant Path Support Configuration With Three Enclosures





Figure 4-21 displays redundant path storage configuration with four enclosures. A single PERC H800 adapter can support up to four disk storage enclosures in a redundant path configuration.


Figure 4-21. Redundant Path Support Configuration With Four Enclosures



 **NOTE:** Ensure that the latest firmware version is installed on your storage controller. You can find the latest firmware and installation instructions on the Dell Support website at support.dell.com.

Reverting From Redundant Path Support to Single Path Support on the PERC H800 Adapter

If you need to revert from redundant path support to single path support, shut down the system and remove the exact same cables that were added to support redundant path support, leaving only one connection between the controller and enclosures. After you remove the cable and power up the system, ensure that there are no warning messages during boot, and that all virtual disks are online and optimal. If you are using Dell OpenManage, see the Dell OpenManage documentation at support.dell.com/manuals for additional instructions.

 **CAUTION:** If you remove any cables other than the ones added to enable redundant path support, the enclosure and disks can get disconnected, and virtual disk may fail.

Driver Installation

The Dell PowerEdge RAID Controller (PERC) H700 and H800 cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the PERC H700 and H800 cards.



NOTE: For more information on VMware ESX drivers, see the VMware ESX documentation at support.dell.com/manuals.



NOTE: To check operating system compatibility, see the Dell Support website at support.dell.com/manuals.

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

- **Installing a driver 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 H700 and PERC H800 family of controllers are already installed and you want to update to the latest drivers.

Installing Windows Drivers

Creating the Driver Media

Perform the following steps to create the driver media:

- 1 Browse to the download section for the system at support.dell.com.
- 2 Locate and download the latest PERC H700 or H800 driver to the system.
- 3 Follow the instructions for extracting the driver to the media.

Pre-Installation Requirements

Before you install the operating system:

- Read the Microsoft *Getting Started* document that shipped with your operating system.
- Ensure that your system has the latest BIOS, firmware, and driver updates. If required, download the latest BIOS, firmware, and driver updates from the Dell Support website at support.dell.com.
- Create a device driver media (diskette, USB drive, CD, or DVD).

Creating the Device Driver Media

Use one of the following two methods to create the device driver media.

Downloading Drivers From the Dell Systems Service and Diagnostic Tools Media

- 1 Insert the *Dell Systems Service and Diagnostics Tools* media in a system. The **Welcome to Dell Service and Diagnostic Utilities** screen is displayed.
- 2 Select your system model and operating system (Microsoft Windows Server 2008).
- 3 Click **Continue**.
- 4 From the list of drivers displayed, select the driver you require. Select the self-extracting zip file and click **Run**. Copy the driver to a diskette drive, CD, DVD, or USB drive. Repeat this step for all the drivers you require.
- 5 During the operating system installation, use the media that you created with the **Load Driver** option to load mass storage drivers. For more information on reinstalling the operating system, see the relevant section for your operating system below.

Downloading Drivers From the Dell Support Website

- 1 Go to support.dell.com.
- 2 Select your line of business.
- 3 Click **Drivers and Downloads**.
- 4 Enter the service tag of your system in the **Choose by Service Tag** field or select your system's model.

- 5 Select the **System Type, Operating System, Driver Language, and Category** from the drop-down list.
- 6 The drivers that are applicable to your selection are displayed. From the available list, download the drivers that you require to a diskette drive, USB drive, CD, or DVD.
- 7 During the operating system installation, use the media that you created with the **Load Driver** option to load mass storage drivers. For more information on reinstalling the operating system, see the relevant section for your operating system below.


Installing Driver During a Windows Server 2003 Operating System Installation

- 1 Boot the system using the Windows Server 2003 media.
- 2 When the message **Press F6 if you need to install a third party SCSI or RAID driver** is displayed in the bottom of the screen, press the <F6> key immediately.


Within a few minutes, a screen is displayed asking for additional controllers in the system.

- 3 Press the <S> key.

The system prompts for the driver media to be inserted.

 **NOTE:** The driver can be provided using a properly formatted USB key. For additional details on the driver, go the Dell Support website at support.dell.com.

- 4 Insert the driver media in the media drive and press <Enter>. A list of PERC H700 and H800 cards is displayed.
- 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 media.

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

Installing Driver During a Windows Server 2008, Windows Server 2008 R2 Installation

- 1 Boot the system using the Windows Vista, Windows Server 2008, Windows 7 Server or Windows Server 2008 R2 media.
- 2 Follow the on-screen instructions until you reach **Where do you want to install Vista/2008/7** window and then select **Load driver**.
- 3 The system prompts you to insert the media. Insert the installation media and browse to the proper location.
- 4 Select the appropriate PERC H700 or H800 card from the list, click **Next** and continue installation.



NOTE: The Windows Server 2008 R2 operating system includes native drivers for the PERC H700 and H800 cards. For driver updates, see the **Drivers and Downloads** section at support.dell.com.

Installing Windows Server 2008, Windows Server 2008 R2, Windows Server 2003 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:

- 1 Turn off the system.
- 2 Install the new RAID controller in the system.
For detailed instructions on installing and cabling the RAID controller in the system, see "Installing and Configuring Hardware" on page 41.
- 3 Turn on the system.
- 4 The **Found New Hardware Wizard** screen displays the detected hardware device.




NOTE: Windows Server 2008 R2 has native drivers for the PERC adapters. The system automatically detects the controller and installs the drivers. Check the version of the driver and update if required.

- 5 Click **Next**.
- 6 On the **Locate device driver** screen, select **Search for a suitable driver for my device** and click **Next**.
- 7 Browse and select the drivers from 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 system when prompted.

Updating Existing Windows Server 2008, Windows Server 2008 R2, Windows Server 2003

 **NOTE:** Close all applications on your system before you update the driver.


- 1 Insert the media (CD, DVD, or USB drive) containing the driver.
- 2 Select **Start**→**Settings**→**Control Panel**→**System**.

The **System Properties** screen is displayed.


 **NOTE:** The path to **System** might vary depending on the operating system family.

- 3 Click on the **Hardware** tab.
- 4 Click **Device Manager**.

The **Device Manager** screen is displayed.

 **NOTE:** The path to **Device Manager** might vary depending on the operating system family.


- 5 Expand **SCSI and RAID Controllers** by double-clicking the entry or by clicking on the plus symbol next to **SCSI and RAID Controller**.

 **NOTE:** In Windows 2008 and Windows Vista, the PERC adapter is listed under **Storage Controllers**.


- 6 Double-click the RAID controller for which you want to update the driver.
- 7 Click the **Driver** tab and click **Update Driver**.

The screen to update the device driver wizard is displayed.


- 8 Select **Install from a list or specific location**.
- 9 Click **Next**.
- 10 Follow the steps in the wizard and browse to the location of the driver files.
- 11 Select the INF file from the driver media (CD, DVD, or other media).
- 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.

 **NOTE:** Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 operating system. DUP is an executable application that updates drivers for specific devices. DUP supports command line interface and silent execution. For more information, see support.dell.com.

Installing Linux Driver

 **NOTE:** PERC H700/H800 cards and both the PERC 5 and PERC 6 family of controllers use the same driver and do not require separate driver installations.

Use the procedures in this section to install the driver for Linux. The driver is updated frequently. To ensure that you have the current version of the driver, download the updated Linux driver from the Dell Support website at support.dell.com.

 **NOTE:** The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below. If not, proceed with using the native device driver and then skip to "Installing the RPM Package With DKMS Support" on page 78.

Creating a DUD

Before beginning the installation, copy the drivers from the *Service and Diagnostic Utilities* media or download the driver appropriate for Linux from support.dell.com. The driver package includes driver RPM (Red Hat Package Managers) file, the DKMS (Dynamic Kernel Module Support) RPM file, driver source code, and release notes.

For more information on DKMS, go to support.dell.com/manuals.

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

- 1 Unzip the package using `gunzip`.
- 2 Untar the file using `tar -xvf`.

The Driver Update Disk (DUD) image can be transferred to a USB flash key, a system floppy disk slot, or a USB floppy device depending upon availability and operating system.

USB key method (Red Hat Enterprise Linux only): Transfer the appropriate .img file to a USB key.

System Floppy Disk slot method: Use the `dd` command to create a driver update disk. Use the appropriate image for the purpose.

- a** Insert a floppy disk into the system floppy disk slot.
- b** At a terminal prompt, do the following:

```
# dd if=<image_file_name> of=/dev/fd0
```

USB Floppy device method: Use the `dd` command to create a driver update disk. Use the appropriate image for the purpose.

- a** Put a floppy disk into a USB floppy device and plug the device into a USB slot of the system under test. Use `dmesg` to find out which device this USB floppy is enumerated to (for example, `sdb`, `sdc`, and so on).
- b** Transfer the driver image to the floppy:

```
# dd if=<image_file_name> of=/dev/sdx
```



NOTE: You can create a driver update disk on a Windows system using the program **dcopynt**.

- 3** Use the diskette for operating system installation. For Red Hat Enterprise Linux, see "Installing Red Hat Enterprise Linux Operating System Using the DUD" on page 76. For SUSE Linux Enterprise Server, see "Installing SUSE Linux Enterprise Server Using the Driver Update Diskette" on page 77.

Creating a Driver Update Diskette Using DKMS

Perform the following steps to create the Driver Update Diskette (DUD) using the DKMS tool:



NOTE: The driver must be installed on the system where this procedure is performed.

- 1 Install the DKMS-enabled `megaraid_sas` driver rpm package.
- 2 Type the following command in any directory:

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



NOTE: The values for the `-d` option are `suse` for Suse Linux Enterprise Server diskettes and `redhat` for RHEL diskettes.



NOTE: For more information on usage of DKMS, see the DKMS main page.

This starts the process to create the `megaraid_sas` DUD image. After the DUD image has been built, you can find it in the DKMS tree for the `megaraid_sas` driver. See the output of the `dkms mkdriverdisk` command for the exact path.

Installing Red Hat Enterprise Linux Operating System Using the DUD

Perform the following steps to install Red Hat Enterprise Linux (versions 4 and 5) and the appropriate driver:

- 1 Boot normally from the Red Hat Enterprise Linux installation media.
- 2 At the command prompt, type: `linux expert dd`
- 3 When the install prompts for additional drivers, insert the diskette or USB key and press `<Enter>`.

For information about creating a driver diskette, see "Creating a DUD" on page 74.

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

Installing SUSE Linux Enterprise Server Using the Driver Update Diskette



NOTE: For information about creating a driver diskette, see "Creating a DUD" on page 74.

To install SUSE Linux Enterprise Server using the DUD:

- 1 Insert the appropriate SUSE Linux Enterprise Server Service Pack media in the system.
- 2 For SUSE Linux Enterprise Server 10, select <F5> for the DUD. For SUSE Linux Enterprise Server 11, select <F6>.

The system displays three options: **Yes**, **No**, and **File**. Choose **Yes** to install the driver.

- 3 Select **Installation** from the menu.
- 4 Press <Enter> to load the Linux kernel.
- 5 At the prompt `Please insert the driver update floppy`, 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 disk and installs it.

Installing the RPM Package With DKMS Support

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

- 1 Uncompress the gzipped tarball driver release package.
- 2 Install the DKMS package using the command: `rpm -ihv dkms-<version>.noarch.rpm`
- 3 Install the driver package using the command: `rpm -ihv megaraid_sas-<version>.noarch.rpm`
 **NOTE:** Use `rpm -Uvh <package name>` when updating an existing package.
- 4 If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
- 5 Verify that the driver has been loaded with the following system commands: `modinfo megaraid_sas` and `dkms status`.

Upgrading the Kernel

When upgrading to a new kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update or install the driver for the new 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 see a message similar to the following one:

```
<driver name>, <driver version>, <new kernel
version>: installed
```
- 3 If the previous device driver is in use, you must reboot the system for the updated driver to take effect.

Installing Solaris Driver

 **NOTE:** The DUD images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below. If not, proceed with the operating system installation using the native device driver and then skip to "Adding or Updating the Driver to an Existing System" on page 80.


Use the procedures in this section to install the driver for Solaris 10. To ensure that you have the current version of the driver, download the updated Solaris driver from the Dell Support website at support.dell.com.


The package is a gzipped .tar file. Download the package to a Solaris system, and perform the following steps:

- 1 Extract the package contents:

```
gunzip -  
c <driver_package.tgz> | tar xvf -
```
- 2 Use the `dd` command to create a driver update disk. Use the appropriate image for the purpose. Type:


```
dd if=./mega_sas.img of=  
/<diskette drive device node> bs=32k
```

 **NOTE:** If you are uncertain which device node corresponds to your diskette drive, execute the `rmformat` command and search for the correct **Logical Node**.

 **NOTE:** You can create a DUD on a system running the Windows operating system using the program `dcopynt`.

- 3 If you prefer, you may use the `cdrecord` command to create a CDROM instead of a floppy image. Type:

```
cdrecord dev=  
<bus>,<target>,<lun> mega_sas_cd.iso
```

 **NOTE:** To identify the correct location of the bus, target, and logical unit number (LUN) combination, execute the following command:

```
cdrecord --scanbus
```

Installing Solaris 10 on a PowerEdge System Booting From a PERC H700 or H800 Card

To install the driver during Solaris 10 operating system installation:

- 1 Boot the system from the Solaris installation media and select the preferred console.
- 2 After Solaris finishes configuring devices, a menu is displayed. Select **Apply Driver Updates**.
- 3 Select [1] if you created a CD from the `mega_sas_cd.iso` file.
- 4 Select [2] if you created a diskette from the `mega_sas.img` file and you are using a traditional diskette drive.
- 5 Select [3] if you created a diskette from the `mega_sas.img` file and you are using a removable (USB) diskette drive.
- 6 Follow instructions presented by **Driver Update** installer.
- 7 The following message is displayed: Installation of <meegasas> was successful.
- 8 Select [e] to end
- 9 Follow the instructions on the screen to complete installation.

Adding or Updating the Driver to an Existing System

- 1 To add the `mega_sas` driver to an existing system, or to upgrade to a newer version of the driver, you must unarchive the driver package and execute the installation script:

```
tar xvf x86_64.tar
cd x86_64
./install.sh
```

- 2 Restart your Solaris system to start using the new driver. To confirm the `mega_sas` driver is loaded, run the following command:

```
modinfo | grep mega_sas
```
- 3 Ensure that the driver version is correct.

Configuring and Managing RAID

Dell OpenManage 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 management applications for PERC H700 and H800 include:

- Dell OpenManage Storage Management
- BIOS Configuration Utility (<Ctrl><R>)



NOTE: Dell Serial Attached SCSI (SAS) RAID Storage Manager is not supported for SED management.

Dell OpenManage Storage Management

The Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring a system's locally-attached RAID and non-RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported 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. Using the Dell OpenManage storage management application, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. The command line interface available on selected operating systems to perform RAID management tasks is fully featured and scriptable. For more information, see the *Dell OpenManage Storage Management User's Guide* at support.dell.com/manuals.

BIOS Configuration Utility

The **BIOS Configuration Utility**, also known as <Ctrl><R>, is a storage management application embedded on the PERC H700 or PERC H800 cards that configures and maintains RAID disk groups and virtual disks. <Ctrl><R> is independent of the operating system.



NOTE: Use the **BIOS Configuration Utility** for initial setup and disaster recovery. You can use advanced features through Dell OpenManage storage management application and Dell SAS RAID storage manager.

The following sections provide information about using the **BIOS Configuration Utility**. For more information, see the online help option by pressing <F1>.



NOTE: The PERC H700 or PERC H800 card configuration utility refreshes the screen to show changes to information. The refresh occurs when you press <F5> or every 15 seconds.

Entering the BIOS Configuration Utility

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

- 1 Turn on the system.

A BIOS screen displays information about the controller and configuration.

- 2 During startup, press <Ctrl><R> when prompted by the BIOS screen.

After you press <Ctrl><R>, if there is only one controller, the **Virtual Disk Management** screen for that controller is displayed. If there are more than one controllers, the main menu screen is displayed. The screen lists the RAID controllers. Use the arrow keys to select the RAID controller you want to configure, and press <Enter> to access the management menus for the controller.



NOTE: You can access multiple controllers through the **BIOS Configuration Utility** by pressing <F12>.



NOTE: You can access PERC 6, and PERC H700 or H800 cards from the same BIOS if the PERC 6 firmware is 6.2.0-0013 or later.

Exiting the Configuration Utility

To exit the **BIOS Configuration Utility**, press <Esc> at any menu screen. If there is only one controller, then a dialog box is displayed to confirm your choice. Select **OK** to exit and press <Enter>.

If multiple controllers are present, then the <Esc> key brings you to the **Controller Selection** screen. Press <Esc> again to reach the exit screen. A dialog box is displayed to confirm your choice. Select **OK** to exit and press <Enter>.

Menu Navigation Controls

Table 6-1 displays the menu keys you can use to move between the different screens in the **BIOS Configuration Utility**.

Table 6-1. Menu Navigation Keys

Notation	Meaning and Use	Example
right-arrow key	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 expands. 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
left-arrow key	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 collapses. Press it again to go to the higher-level menu. Word wrap is supported.	Controller 0 ← Disk Group 1
up-arrow key	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

Table 6-1. Menu Navigation Keys (continued)

Notation	Meaning and Use	Example
down-arrow key	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 Disk 1 ↓ Virtual Disk 4
<Enter>	After you highlight a menu item, press <Enter> to select that item. An options menu for the menu item opens. It 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.	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.
<Tab>	Press <Tab> to move the cursor to the next control on a dialog box 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 Sort By to the previously selected PD in the PD Mgmt screen
<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 .	Press <Ctrl><N> on the VD Mgmt screen to move to the PD Mgmt screen.

Table 6-1. Menu Navigation Keys (continued)

Notation	Meaning and Use	Example
<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 .	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>
<F5>	Press <F5> to refresh the information on the screen.	<F5>
<F11>	Switch between two controllers.	<F11>
<F12>	Press <F12> to display a list of controllers.	<F12>
Spacebar	Press the spacebar to select an item.	Press the <spacebar> to select or deselect a controller setting in the Ctrl Mgmt View .

Setting Up Virtual Disks

You can set up a disk group and create virtual disks using the procedures contained in this section. 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).

For more information, see "Managing Dedicated Hot Spares" on page 97.

- 3 Initialize the virtual disks.



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

When you define the virtual disks, you can set the following virtual disk parameters:

- RAID level
- Stripe element size
- Read policy
- Write policy
- Type of initialization
- Hot spare configuration



NOTE: The default hard drive cache policy for a virtual disk composed with SAS hard drives is *disabled* and with SATA hard drives is *enabled*. The Virtual Disk parameter can not be changed in the **BIOS configuration Utility**.

Table 6-2 shows the parameters that you can configure when defining virtual disks.

Table 6-2. Virtual Disk Parameters and Descriptions


Parameter	Description
RAID Level	RAID Level specifies whether the virtual disk is RAID 0, 1, 5, 6, 10, 50, or 60. The number of disks, disk capacity, the requirements for fault tolerance, performance, and capacity should be considered when selecting the RAID level. For more information, see "Summary of RAID Levels" on page 17.
Stripe Element Size	Stripe Element Size specifies the size of the segments written to each physical disk in a RAID 0, 1, 5, 6, 10, 50, and 60 virtual disk. You can set the stripe element size to 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB. The default and recommended stripe element size is 64 KB. A larger stripe element size provides better read performance if your system mostly does sequential reads.


Table 6-2. Virtual Disk Parameters and Descriptions (continued)

Parameter	Description
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: If a Battery Backup Unit (BBU) is present, the default cache setting is Write-Back. If no BBU is present, the default cache policy default setting is Write-Through.</p> <p>NOTE: 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 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>
Read Policy	<p>Read-Ahead enables the read ahead feature for the virtual disk. You can set the parameter to Read-Ahead, No-Read-Ahead, or Adaptive. The default is Adaptive-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 is required soon.</p> <p>No-Read-Ahead specifies that the controller does not use read ahead for the current virtual disk.</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>

Virtual Disk Management

Creating Virtual Disks

 **NOTE:** Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and SSDs within a virtual disk is not supported.

 **NOTE:** To create secured virtual disks, see "Security Key and RAID Management" on page 121.

Perform the following steps to create a virtual disk:

- 1 During host system bootup, press <Ctrl><R> when the BIOS screen is displayed.

The **Virtual Disk Management** screen is displayed. If there is more than one controller, the main menu screen is displayed. Select a controller, and press <Enter>. The **Virtual Disk Management** screen is displayed for the selected controller.

- 2 Use the arrow keys to highlight **Controller #** or **Disk Group #**.

- 3 Press <F2>

The list of available actions is displayed.

- 4 Select **Create New VD** and press <Enter>.

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

When adding a virtual disk to a Disk Group, the **Add VD in Disk Group** screen is displayed. Skip to step 11 to change the basic settings of the virtual disk.

- 5 Press <Enter> to display the possible RAID levels, based on the physical disks available.

- 6 Press the down-arrow key to select a RAID level and press <Enter>.


- 7 When creating a spanned virtual disk (RAID 10, 50 or 60), enter the number of physical disks per span in the **PD per Span** field and press <Enter>.

 **NOTE:** Creating a 22 physical disk RAID 10 virtual disk is possible by selecting RAID 10 and populating the **PD per Span** field with 22.

- 8 Press <Tab> to move the cursor to the list of physical disks.

- 9 Use the arrow key to highlight a physical disk and press the spacebar, <Alt>, or <Enter> to select the disk.
- 10 Select additional disks, if required.
- 11 Press <Tab> to move the cursor to the **Basic Settings** box.
- 12 Set the virtual disk size in the **VD Size** field.

The virtual disk size is displayed in GB format.

 **NOTE:** 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.

- 13 Press <Tab> to access the **VD Size** field, and type a virtual disk name.
- 14 Press <Tab> to move the cursor to **Advanced Settings**.
- 15 Press the spacebar to make the settings active so that you can change them.

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

The defaults parameters are displayed. You can accept the defaults or change them. For detailed information about the virtual disk parameters, see "Virtual Disk Parameters and Descriptions" on page 86.


- 16 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, 128 KB, 256 KB, 512 KB, and 1024 KB). Press the down-arrow key to highlight the desired option and press <Enter>. The default is 64 KB.
 - e If you need to change the read policy, press <Tab> to move the cursor to **Read Policy**.

- f Press <Enter> to display the options (**No Read Ahead**, **Read Ahead**, or **Adaptive Read Ahead**). Press the down-arrow key to highlight the desired option and press <Enter>.
- g If you need to change the write policy, press <Tab> to move the cursor to **Write Policy**.
- h Press <Enter> to display the options (**Write-Through**, **Write Back**). Press the down-arrow key to highlight an option and press <Enter>.
- i Press <Tab> to move the cursor to **Force WB with no battery** and press <Enter>. If you chose **Write-Through** as the write policy, then the option is not available.
- j Press <Tab> to move the cursor to **Initialize** and press <Enter>.

 **CAUTION: Do not initialize virtual disks when attempting to recreate an existing configuration.**



NOTE: The initialization performed at the stage is fast initialization.

- k Press <Tab> to move the cursor to **Configure HotSpare** and press <Enter>.
-  **NOTE:** The hot spare created at the stage is a dedicated hot spare.
- l If you have chosen to create hot spares in the earlier steps, a pop-up window is displayed where disks with appropriate sizes are displayed. Press the spacebar to select the disk size.
- m After you select the disk size, click **OK** to finalize the selection or click **Cancel** to forfeit the selection.
- n Select **OK** to accept the settings and press <Enter> to exit the window or select **Cancel** and press <Enter> to exit if you do not want to change any virtual disk parameters.

Initializing Virtual Disks



CAUTION: A Full Initialization permanently destroys any existing data on that virtual disk.

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.
- 4 A pop-up windows is displayed indicating that the virtual disk has been initialized.
- 5 Repeat the procedures from step 1 to step 4 to configure another virtual disk.

The PERC H700 or PERC H800 cards support up to 64 virtual disks per controller. The currently configured virtual disks display on the screen.

Checking Data Consistency

Select the **Consistency Check (CC)** option in the configuration utility to verify the redundancy data in virtual disks that use RAID levels 1, 5, 6, 10, 50, and 60. (RAID 0 does not provide data redundancy.)

If you attempt to run a **Consistency Check** on a virtual disk that has not been initialized, the following error message is displayed:

```
The virtual disk has not been initialized. Running a consistency check may result in inconsistent messages in the log. Are you sure you want to continue?
```

You can select **Yes** or **No**. If you select **Yes**, the CC operation continues. If you select **No**, the operation ends.

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 available actions (**Start, 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 if needed.

Importing or Clearing Foreign Configurations Using the VD Mgmt Menu

When a foreign configuration exists, the BIOS screen displays the message `Foreign configuration(s) found on adapter`. In addition, a foreign configuration is displayed on the right side of the **Ctrl Mgmt** screen.

You can use the **VD Mgmt** menu to import the existing configuration to the RAID controller or clear the existing configuration. In addition, you can view the foreign configuration from the **Foreign View** tab without importing the configuration.



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



NOTE: To import a secured foreign configuration, see "Security Key and RAID Management" on page 121.

Perform the following steps to import or clear foreign configurations:

- 1 During bootup, press <Ctrl> <R> when prompted by the BIOS screen.
The **VD Mgmt** screen is displayed by default.
- 2 On the **VD Mgmt** screen, highlight the **Controller #**.
- 3 Press <F2> to display the available actions.

- 4 Navigate to the **Foreign Config** option and 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** screen displays detailed configuration information. It 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



NOTE: To import a secured foreign configuration, see "Security Key and RAID Management" on page 121.

If one or more physical disks are removed from a configuration, the configuration on those disks is considered a foreign configuration by the RAID controller.

You can use the **Foreign Configuration View** screen to view information about the foreign configuration, such as disk groups, virtual disks, physical disks, space allocation, and hot spares. The foreign configuration data is displayed in the same format as configurations on the **VD Mgmt** screen. You can use the **VD Mgmt** screen to view the foreign configuration before importing. After you view the foreign configuration, you can either clear or import to the RAID controller.



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

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.

The following constraints apply to the physical disks that are considered for import:

- The disk state of a physical disk can change from the time the foreign configuration is scanned to when the actual import occurs. The foreign import occurs only on disks that are in the **Unconfigured Good** state.
- Disks in the failed or offline state cannot be imported.
- The firmware does not allow you to import more than eight foreign configurations.


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 reinserted, the controller considers the disks 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**, **Clear**).

 **NOTE:** You must have all the disks in the system 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).

In the **Preview Configuration Data** window, the status of a physical disk that needs to be rebuilt is displayed as **Rebuild**.

 **NOTE:** When you import a foreign configuration, the dedicated hot spares in the configuration are imported as dedicated hot spares on two conditions — the associated virtual disk is already present or the associated virtual disk is also imported along with the configuration.

 **NOTE:** Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. For more information about checking data consistency, see "Checking Data Consistency" on page 91.

2 If all the physical disks in a virtual disk are removed at different times and re-inserted, the controller considers the disks 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** and **Clear**.



NOTE: You must have all the drives in the system 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 are imported, and then automatically rebuilt.



NOTE: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. For more information about checking data consistency, see "Checking Data Consistency" on page 91.

3 If the physical disks in a non-redundant virtual disk are removed, the controller considers the disks 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** and **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 occur after the import operation because there is no redundant data to rebuild the disks with.

Managing Preserved Cache

If a virtual disk becomes offline or is deleted because of missing physical disks, the controller preserves the dirty cache from the virtual disk.

The preserved dirty cache, known as pinned cache, is preserved until you import the virtual disk or discard the cache.



NOTE: Certain operations, such as creating a new virtual disk, cannot be performed if preserved cache exists. You have to enter the **BIOS Configuration Utility** to resolve the situation before you boot to the operating system. Messages are displayed notifying you that you must enter the **BIOS Configuration Utility** to discard the preserved cache or import the virtual disks with the preserved cache.



CAUTION: If there are any foreign configurations, it is strongly advised that you import the foreign configuration before you discard the preserved cache. Otherwise, you might lose data that belongs with the foreign configuration.

Perform the following steps to select whether to import the virtual disk or discard the preserved cache:

- 1 On the **VD Mgmt** screen, click on a controller icon.
- 2 Press <F2> to display the menu of available actions.
- 3 Select **Manage Preserved Cache**.

A message is displayed advising you to import the foreign configuration before you discard the preserved cache to avoid losing data belonging to the foreign configuration. Confirm whether you want to continue. The **Manage Preserved Cache** screen displays the affected virtual disks.

- 4 You can choose to discard the cache on the **Manage Preserved Cache** screen. If you press **Cancel**, the process is cancelled and the **Preserved Cache Retained** dialog box is displayed.

If you choose to discard the cache, you are prompted to confirm your choice. If you choose to retain the cache, a message is displayed to notify you that you cannot perform certain operations while the cache exists. Click **OK** to continue.

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 a global hot spare 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 with an **X** beside them and the physical disks that are available to create dedicated hot spares.

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

- 3 Use the following instructions to create or delete a dedicated hot spare:

- **Creating a dedicated hot spare**

- a Press the down-arrow key to highlight an available physical disk.
- b Press the spacebar to select the disk.
- c Repeat step a to step b for each dedicated hot spare that you want to create.


An **X** is displayed beside the selected physical disk(s).

- **Deleting a dedicated hot spare**

- a Use the down-arrow key to highlight a current hot spare.
- b Press the spacebar to de-select the disk.
- c Repeat step a to step b for each dedicated hot spare that you want to delete.

- 4 Press <Enter> to confirm the changes.

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

 **NOTE:** If a global hot spare or dedicated hot spare is removed, reinserted and imported, it regains its status as a hot spare. A dedicated hot spare becomes a global hot spare when the disk group it was assigned to protect is no longer present during import.

Deleting Virtual Disks



NOTE: You cannot delete a virtual disk during an initialization.



NOTE: Warning messages appear stating the effect of deleting a virtual disk. Click **OK** twice to complete the virtual disk deletion.

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

- 1 Press <Ctrl><N> to access the **VD Mgmt** screen.
- 2 Use the arrow keys to move the cursor to the **Virtual Disks** heading.
- 3 Press <F2>.
The action menu is displayed.
- 4 Select **Delete VD** and press <Enter>.
- 5 If there are multiple virtual disks in a **Disk Group**, select **Total Free Capacity** for the Disk Group in the **VD Mgmt** screen.
The total amount of free space available in the **Disk Group** is displayed.

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 Use the arrow keys to move the cursor to the **Virtual Disks** heading.
- 3 Press <F2>.
The action menu is displayed.
- 4 Select **Delete Disk Group** and press <Enter>.
The disk group is deleted.

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

Clearing the Configuration

You can delete all virtual disks on the RAID controller by performing the operation.

To clear the configuration, perform the following steps in the **BIOS Configuration Utility**:

- 1 Press <Ctrl><N> to access the **VD Mgmt** screen.
- 2 Use the arrow keys to move the cursor to the **Controller** heading.
- 3 Press <F2>. The action menu is displayed.
- 4 Select **Clear Config**.
A pop-up window is displayed prompting for confirmation to delete all virtual disks.
- 5 Select **NO** to delete the virtual disks or **YES** to retain the existing configuration.

BIOS Configuration Utility Menu Options

The first menu that is displayed when you access the **BIOS Configuration Utility** is the main menu screen. It lists the controller, controller number, and other information, such as the slot number. On the screen, you can use the arrow keys to select the RAID controller you want to configure. Press <Enter> to access the controller.

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 (Foreign View)** menu

Most menus consist of two panels:

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

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 is displayed when you access a RAID controller from the main menu screen on the **BIOS Configuration Utility**. The left panel displays the menus for the virtual disk management, which are:

- **Controller #**
 - **Disk Group #**
 - **Virtual Disks**
 - **Physical Disks**
 - **Total Free Capacity** (virtual disk size and free space you can use to create a virtual disk)
 - **Hot Spares** (global and dedicated)

The right panel displays detailed information for the selected controllers, disk groups, virtual disks, physical disks, total free capacity, and hot spares, as shown in Table 6-3.

Table 6-3. Information on the Virtual Disk Management Screen

Menu Item Selected in Left Panel	Information Displayed in Right Panel
Controller	Controller Properties: <ul style="list-style-type: none">• Number of disk groups (DG)• Number of virtuals disks (VD)• Number of physical disks (PD)

Table 6-3. Information on the Virtual Disk Management Screen (continued)

Menu Item Selected in Left Panel	Information Displayed in Right Panel
Disk Group #	Disk Group # Properties: <ul style="list-style-type: none">• Number of virtual disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares• Security property of the Disk Group
Virtual Disks	Disk Group # Properties: <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 #	Virtual Disk # Properties: <ul style="list-style-type: none">• RAID level (0, 1, 5, 6, 10, 50, or 60)• RAID status of the virtual disk (Failed, Degraded, Partially Degraded, or Optimal)• Operation currently in progress Disk Group # Properties: <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

Table 6-3. Information on the Virtual Disk Management Screen (continued)

Menu Item Selected in Left Panel	Information Displayed in Right Panel
Physical Disks	Disk Group # Properties: <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 #	Physical Disk Properties: <ul style="list-style-type: none">• Vendor name• Physical disk state• Enclosure Position• Slot Position Disk Group # Properties: <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
Total Free Capacity	Disk Group # Properties: <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

Table 6-3. Information on the Virtual Disk Management Screen (continued)

Menu Item Selected in Left Panel	Information Displayed in Right Panel
Hot Spares	Physical disk properties: <ul style="list-style-type: none">• Vendor name• Physical disk size• Physical disk state• Enclosure Position• Slot Position Disk group # properties: <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 Disk Actions

Table 6-4 describes the actions you can perform on virtual disks. For procedures you can use to perform the actions, See "Virtual Disk Management" on page 88.

Table 6-4. Virtual Disk Actions

Action	Description
Create a new virtual disk	Creates a new virtual disk from one or more physical disks. You can configure hot spares when you create a virtual disk.
Manage dedicated hot spares	Creates or deletes a hot spare that you can dedicate to a single redundant virtual disks.
Initialize a virtual disk	Initializes the selected virtual disk. You must initialize every virtual disk that is configured. You can perform a fast initialization or a Full Initialization.
Check data consistency on a virtual disk	Verifies the correctness of the redundancy data in the selected virtual disk. The option is available only if RAID level 1, 5, 6, 10, 50, or 60 is used. The PERC H700 or PERC H800 cards 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 and read policy from the menu.
Manage preserved cache	Preserves the dirty cache from a virtual disk if it becomes offline or is deleted. The dirty cache is preserved until you import the virtual disk or discard the cache.
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 the headings. You can perform several actions on the physical disks, including the following:

- Rebuilding physical disks
- Performing the Replace Member operation
- Setting the LED to blink
- Making a disk online or offline (unaffiliated with a disk group)
- Creating global hot spares
- Removing dedicated hot spares or global hot spares

The **PD Mgmt** screen also displays several physical disk properties as shown in Table 6-5.

Table 6-5. Information on the Physical Disk Management Screen

Information Displayed in Left Panel	Supported Information Displayed in Right Panel
Physical Disk: <ul style="list-style-type: none">• Disk ID• Protocol type• Capacity (GB)• Physical Disk State• Disk Group• Vendor	<ul style="list-style-type: none">• Security Property of Physical Disk• Encryption Capable• Product ID• Firmware Revision• Disk Write Cache• S.M.A.R.T state• Physical Disk operation• Max Device Link Rate• Negotiated Link Rate• Dell Certified Disk

Physical Disk Actions

Table 6-6 describes the actions you can perform on physical disks. For procedures that can be used to perform the actions, see "Physical Disk Management" on page 108.

Table 6-6. Physical Disk Actions

Action	Description
Rebuild	Regenerates all data to a replacement disk in a redundant virtual disk (RAID level 1, 5, 6, 10, 50, or 60) after a disk failure. A disk rebuild normally occurs without interrupting normal operations on the affected virtual disk.
Replace Member	Replaces the disk in the virtual disk with another disk that can be selected.
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. Designates the selected physical disk as a global hot spare.
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. For information on performing a physical disk rebuild, see "Performing a Manual Rebuild of an Individual Physical Disk" on page 111.

Several of the controller configuration settings and the virtual disk settings affect the actual rate of rebuild. The factors include the rebuild rate setting, virtual disk stripe size, virtual disk read policy, virtual disk write policy, and the amount of workload placed on the storage subsystem. For information on getting the best rebuild performance from your RAID controller, see the documentation on Dell Support website at support.dell.com/manuals.

Controller Management (Ctrl Mgmt)

The **Controller Management** screen (**Ctrl Mgmt**) displays the product name, package, firmware version, BIOS version, boot block version, controller ID, security capability, and security key presence. Use the screen to perform actions on the controller and BIOS. You can perform functions such as enable or disable the controller BIOS, enable or disable the BIOS during bootup in the event of BIOS errors, and enable or disable the option to **Auto Import**. In addition, you can select a virtual disk from which to boot, and select default settings.

Controller Management Actions

Table 6-7 describes the actions you can perform on the **Ctrl Mgmt** screen.

Table 6-7. Controller Management Options

Option	Description
Enable Controller BIOS	Select the 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 BIOS Stop On Error	Select the option to stop the system BIOS during bootup if there are BIOS errors. The option enables you to enter the configuration utility to resolve the problem.
Select Bootable virtual disk	Select the option to specify a virtual disk as the boot disk on the controller. The option is displayed if you have built virtual disks.
Enable Auto Import	Attempts to import every online foreign configuration during boot without having the need to access the BIOS Configuration Utility .
Factory Default	Select the option to restore the default settings for the options in the Settings box.

Foreign Configuration View

When a foreign configuration is present, you can select **Foreign Configuration View** to display the configuration. The screen shows the foreign configuration as it would be if you import it. You can preview the foreign configuration before you decide whether to import it or clear it.

In some cases, a foreign configuration cannot be imported. If a physical disk in a virtual disk is rebuilding, the physical disk's state is set to **Rebuild**. No virtual disk target ID displays for virtual disks that cannot be imported.

The section "Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen" on page 93 contains the procedures you can use to manage the foreign configurations.



NOTE: The **BIOS Configuration Utility** reports error codes for failed imports of foreign configurations.

Physical Disk Management

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 the option:

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.
A list of physical disks is displayed. The status of the each disk is displayed 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.

Creating Global Hot Spares

You can use a global hot spare 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 is displayed. The status of the each disk is displayed 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 is displayed under the heading **State**.

 **NOTE:** To replace a failed physical disk global hot spares must use the same disk 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 is displayed. The status of each disk is displayed 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 Hot Spare** from the list of actions and press <Enter>.

The physical disk is changed to the **Ready** state. The status of the physical disk is displayed 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 step 1 to step 4 to remove them.

Replacing an Online Physical Disk

In addition to the automatic **Replace Member** operation, you can manually replace any physical disk that is part of a virtual disk using the **Replace Member** functionality. Perform the following steps to replace a physical disk:

- 1 In the **Virtual Disk Management** window, select **Virtual Disk #** and press the down-arrow key until **Physical Disks** is highlighted.
- 2 Press the right-arrow key to expand the list of physical disks that are members of the virtual disk.
- 3 Press the down-arrow key and highlight the desired physical disk you want to replace. Press <F2> to expand the list of allowed operations on the disk.
- 4 Select **Replace** and then **Start**.
- 5 Press the down-arrow to highlight a replacement disk and then press the spacebar to select the disk.
- 6 Select **OK** to start the replacement.



NOTE: The replacement disk must be a hot spare or an unconfigured disk without a foreign configuration. It must have the same or greater capacity and should be of the same type as the disk it is replacing.

Restrictions and Limitations

The following restrictions and limitations apply to the **Replace Member** operation:

- The **Replace Member** functions are restricted to one per array for RAID 0, RAID 1, and RAID 5, and two per array for RAID 6.
- The **Replace Member** function and rebuild cannot run simultaneously on a RAID 6 virtual disk. The rebuild operation has a higher priority, and the **Replace Member** operation is aborted if a rebuild begins.
- The **Replace Member** function cannot replace a SED with a non-SED if the virtual disk is secured.

Stopping Background Initialization

Background initialization (BGI) is the automated operation in which parity is created and written. BGI does not run on RAID 0 virtual disks. Under certain conditions, the **BIOS Configuration Utility** displays a message if you want to stop BGI in progress. An alert message is displayed if BGI is in progress and you start any of the following actions:

- A Full Initialization on the virtual disk
- A Fast Initialization on the virtual disk
- A Consistency Check on the virtual disk

The following alert message is displayed: 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 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 is displayed. The status of each disk is displayed 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.

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

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



NOTE: You can also use the **VD Mgmt** screen to perform a manual rebuild. Use the arrow key to highlight a physical disk, and press <F2>. In the menu that is displayed, select the **Rebuild** option.



CAUTION: If a physical disk is a member of a disk group that contains multiple virtual disks and one of the virtual disks is deleted when a rebuild operation is in progress, the rebuild operation stops. You can then resume the rebuild operation manually using a storage management application. To avoid interruption, ensure that none of the virtual disks are deleted until the rebuild is complete.

Controller Management

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, you can enable BIOS on multiple controllers. However, if you want to boot from a specific controller, 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** is displayed beside **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 de-select 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.
- 3 Press the down-arrow key to display a list of virtual disks.
- 4 Use the down-arrow key to highlight a virtual disk.
- 5 Press <Enter> to select the 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** is displayed beside **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 de-select **Enable BIOS Stop on Error**, then select **Apply** and press <Enter>.

Enabling Auto Import

If there is a native configuration present on the controller, the option **Enable Auto Import** automatically imports every online foreign configuration during boot without having the need to access the **BIOS Configuration Utility**.



NOTE: The controller automatically imports every optimal and degraded foreign configuration without enabling the feature if there is no native configuration on the controller.

To enable Auto Import:

- 1 Press <Ctrl><N> to access the **Ctrl Mgmt** menu screen.
- 2 Press <Tab> to move the cursor to **Enable Auto Import** in the **Settings** box.
- 3 Press the spacebar to select **Enable Auto Import**.
An X is displayed beside Enable Auto Import.
- 4 Press <Tab> to move the cursor to the **Apply** button, and then press <Enter> to apply the selection.
The **Auto Import** is enabled.

To disable Auto Import:

- 1 Use the spacebar to de-select **Enable Auto Import**.
- 2 Select **Apply** and press <Enter>.
The **Auto Import** is disabled.

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 BIOS Stop on Error**, and **Enable Auto Import**. 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 de-select the settings for the options in the **Settings** box.
- 4 Press <Tab> to move the cursor to the **Factory Default** box, and press the <Alt>, <Enter>, or the spacebar.
A dialog box is displayed for you to confirm your choice.
- 5 Select <OK> and press <Enter>.

The defaults are automatically selected for the controller settings and are displayed in **Settings**.

CacheCade

The Dell PowerEdge RAID Controller (PERC) H700 and H800 cards support CacheCade, a feature that can improve application performance by increasing read caching capacity. The CacheCade feature makes use of high-performing solid state disks (SSDs) as a secondary tier of cache. CacheCade provides faster reads and maximizes transactional I/O performance.

The use of SSDs for caching allows a large quantity of data to be present in the cache, resulting in performance improvement in read-intensive applications. Some examples of read-intensive applications include online transaction processing (OLTP), file server, and web server workloads. CacheCade allows for an increase in the I/O performance of hard disk drive (HDD)-based disk groups with the assistance of SSD technology.

CacheCade Virtual Disk Characteristics

The CacheCade feature has the following characteristics:

- Support for CacheCade virtual disks exists only on controllers containing 1 GB of Non-Volatile (NV) Cache.
- CacheCade virtual disks can only be created with SSDs.
- The maximum combined size of CacheCade virtual disks is 512 GB.
 - ✍ **NOTE:** Multiple CacheCade virtual disks may be created, but they are combined to operate as a single cache pool up to the maximum size.
- Data on virtual disks containing secured Self-Encrypting Disks (SEDs) or SSDs will not be cached by CacheCade.
 - ✍ **NOTE:** Data on VDs with unsecured SEDs can be cached using CacheCade.
- CacheCade virtual disks will only cache input reads that are smaller than 64 KB.
- CacheCade virtual disks are read cache only.
- CacheCade virtual disks will not migrate to a controller that does not support CacheCade.

- Importing a CacheCade drive may result in a RAID 0 VD. You will need to reconfigure the CacheCade VD after importing.
- CacheCade virtual disks are not presented to the operating system.

Configuring and Managing CacheCade Virtual Disks

The Dell OpenManage storage management application and the controller's BIOS Configuration Utility (<Ctrl> <R>) allow the creation and deletion of CacheCade virtual disks.

The following sections describe the menu options specific to CacheCade virtual disk management and provide detailed instructions to perform the configuration tasks. The contents of the following section apply to the BIOS Configuration Utility. For more information on the management applications, see "Configuring and Managing RAID" on page 81.

CacheCade Virtual Disk Management

The **Virtual Disk Management** screen is the first screen that is displayed when you access a RAID controller from the main menu screen on the BIOS Configuration Utility.

The following are CacheCade-related actions you can perform through the virtual disk management menu:

- Create CacheCade virtual disk
- Assign CacheCade virtual disk name
- Delete virtual disk

Create CacheCade Virtual Disk



NOTE: Only SSDs can be used to create CacheCade virtual disks.



NOTE: Combining SAS and SATA SSDs within a CacheCade virtual disk is not supported.



NOTE: To avoid inefficient use of space, it is recommended that CacheCade volumes are created with drives of the same size. Larger disk sizes are truncated to the size of the smallest contributing disk, similar to RAID 0.

To create a CacheCade virtual disk:

- 1 During host system bootup, press <Ctrl><R> when the BIOS screen is displayed.

The Virtual Disk Management screen is displayed. If there are more than one controller, the main menu screen is displayed.

- 2 Select a controller, and press <Enter>.

The Virtual Disk Management screen is displayed for the selected controller.

- 3 Use the arrow keys to highlight **Controller #**.


- 4 Press <F2>.

The list of available actions is displayed.

- 5 Select **Create CacheCade Virtual Disk** and press <Enter>.

The **Create CacheCade Virtual Disk** screen is displayed. The cursor is on the first SSD listed in the **Select SSD** section.


- 6 Select the desired SSD(s). As each new SSD is selected, the **CacheCade Virtual Disk Size** changes to reflect the new size.

 **NOTE:** You cannot specify the size of the virtual disk.

- 7 Press <Tab> to move the cursor to the **CacheCade Virtual Disk Name** field. Enter a name if required.

- 8 After you specify the virtual disk name, select **OK** to save the selection or select **Cancel** to cancel the selection.


After the CacheCade virtual disk is created successfully, it is listed in the **Virtual Disk Management** screen under the CacheCade disk group, and is labeled as a CacheCade virtual disk. The virtual disk has an optimal state and its RAID level is RAID 0.


 **NOTE:** You can only delete or rename a CacheCade virtual disk. Background Initialization, fast initialization, full initialization, and consistency check operations are not applicable to CacheCade virtual disks.

Delete CacheCade Virtual Disk

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

- 1 Press <Ctrl><N> to access the **Virtual Disk Management** screen.
- 2 Use the arrow keys to move the cursor to the **CacheCade Disk Group** or **Virtual Disks** heading.
- 3 Press <F2>.
The **Action** menu is displayed.
- 4 Select **Delete VD** and press <Enter>.

 **NOTE:** Warning messages are displayed stating the effect of deleting a virtual disk. Click **OK** to complete the virtual disk deletion.

 **NOTE:** In operating system management applications, a CacheCade virtual disk deletion or removal can be done without interrupting any I/Os. The controller stops caching via the secondary cache, but all outstanding I/Os are completed.

Reconfiguring CacheCade Virtual Disks

In operating system management applications, the resizing of CacheCade virtual disks occurs without stopping any existing I/Os. Creating a new CacheCade virtual disk or adding one or more SSDs to an existing CacheCade virtual disk increases the total cache size. The new resources are immediately used after the addition.

There are two methods to reconfigure CacheCade virtual disks:

- Automatic reconfiguration of cachecade virtual disks
- Manual reconfiguration of cachecade virtual disks

Automatic Reconfiguration of CacheCade Virtual Disks

A CacheCade virtual disk that is made up of more than one SSD is automatically reconfigured upon a removal or failure of a member SSD. The virtual disk retains an Optimal state and adjusts its size to reflect the remaining number of member disks. If auto-rebuild is enabled on the controller, when a previously removed SSD is inserted back into the system or replaced with a new compatible SSD, the CacheCade automatically

reconfigures and adjusts its size to reflect the addition of the member SSD. The number of SSDs to be removed from a CacheCade virtual disk cannot equal the total number of SSDs currently in the CacheCade virtual disk.

After the automatic reconfiguration and resizing of a CacheCade virtual disk, the new virtual disk size is displayed in both the BIOS configuration utility as well as in the OpenManage storage management application.



CAUTION: If a disk is removed from a CacheCade virtual disk, the associated slot for that disk is a hot slot for the CacheCade volume. Any compatible disk inserted into that slot is automatically added to the CacheCade virtual disk. Any pre-existing data on that drive is lost.



NOTE: Disks inserted into a CacheCade hot slot must be equal to or greater than the smallest contributing disk in the virtual disk.

Manual Resizing of CacheCade Virtual Disks

In operating system management applications, a manual resizing of a CacheCade virtual disk occurs due to the addition or removal of drives. Reboot is not necessary.

- Any number of SSDs can be added to a CacheCade virtual disk.
- There is no SAS and SATA mixing allowed within a CacheCade virtual disk, so SATA SSDs cannot be added to a SAS CacheCade virtual disk and vice versa.
- HDDs cannot be added to a CacheCade virtual disk.



NOTE: Capacity of drives added to a CacheCade volume must be equal to or greater than the smallest contributing drive in the virtual disk. The manual resizing of a CacheCade virtual disk cannot be initiated in the BIOS configuration utility. It can only be initiated in the OpenManage storage management application.

Security Key and RAID Management

Security Key Implementation

Dell PowerEdge RAID Controller (PERC) H700 and H800 cards support Self-Encrypting Disks (SED) for protection of data against loss or theft of SEDs. Protection is achieved by the use of encryption technology on the drives. The encryption key is protected from unauthorized use by a security key.

There is one security key per controller. Under Local Key Management (LKM) the key is managed by you (on-controller key management). The key can be escrowed in to a file using Dell OpenManage. The security key is used by the controller to lock and unlock access to encryption-capable physical disks. In order to take advantage of this feature, you must:

- 1 Have SEDs in your system.
- 2 Create (LKM) a security key.

Configuring and Managing Secured Virtual Disks

The Dell OpenManage storage management application and the controller's BIOS Configuration Utility (<Ctrl><R>) allow security keys to be created and managed as well as the creation of secured virtual disks.

The following sections describe the menu options specific to security key management and provide detailed instructions to perform the configuration tasks. The contents of the following section apply to the BIOS Configuration Utility. For more information on the management applications, see "Configuring and Managing RAID" on page 81.

BIOS Configuration Utility Security Menu Options

The BIOS Configuration Utility is a storage management application that resides in the controller BIOS. Its operation is independent of the operating systems. It allows you to configure and maintain physical disk groups and virtual disks as well as provides security key management.

Virtual Disk Management (VD Mgmt)

The Virtual Disk Management screen, **VD Mgmt**, is the first screen that is displayed when you access a RAID controller from the main menu screen on the BIOS Configuration Utility.

Virtual Disk Security Actions

The following are security-related actions you can perform through the virtual disk management menu:

- Security Key Management: creates, changes, or deletes the security settings on a controller.
- Secure Disk Group: secures all virtual disks in Disk Group.

Physical Disk Management (PD Mgmt)

The Physical Disk Management screen displays physical disk information and action menus.

Physical Disk Security Actions

The following are security-related actions you can perform through the physical disk management menu:

- Instant Secure Erase: Resets the security attributes of the SED, rendering existing data inaccessible.

For more information on the Physical Disk Management screen, see "Physical Disk Management" on page 108.

Security Key Management

Local Key Management (LKM)

In local key management (LKM) you generate the key ID and the passphrase required to secure the virtual disk. You can secure virtual disks, change security keys and manage secured foreign configurations using this security mode.



NOTE: Under LKM, you are prompted for a passphrase when you create the key.

Create Key

Perform the following steps to create a local security key on the controller:



NOTE: There is no passphrase backup option when you create a security key; you need to remember your passphrase.

- 1 During the host system boot up, press <Ctrl><R> when the BIOS screen is displayed.

The **Virtual Disk Management** screen is displayed.

If there is more than one controller, the main menu screen is displayed.

- 2 Select a controller, and press <Enter>.

The **Virtual Disk Management** screen is displayed for the selected controller.

- 3 Use the arrow keys to highlight **Security Key Management**.

- 4 Press <F2> to display the actions you can perform.

- 5 Select **Create Key** and press <Enter>.

The **Create Security Key** screen is displayed. The cursor is at the **Security Key Identifier**.

- 6 Enter an identifier for your security key.



NOTE: The **Security Key Identifier** is a user-supplied clear text label used to associate the correct security key with the controller.


- 7 Press <Tab> to enter a passphrase.

- 8 Press <Tab> and select **OK** to accept the settings and to exit the window. Select **Cancel** to exit if you do not want to create a security key on the controller.




CAUTION: If you lose or forget your passphrase, the data in your virtual disk is inaccessible.

Change Key


 **NOTE:** Change Key is active if there is a security key present on the controller.

Perform the following steps when changing the security key on the controller:

- 1 During host system bootup, press <Ctrl><R> when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
- 2 Select a controller, and press <Enter>.
The **Virtual Disk Management** screen is displayed for the selected controller.
- 3 Use the arrow keys to highlight **Security Key Management**.
- 4 Press <F2> to display the actions you can perform.
- 5 Select **Change Key** and press <Enter>.
- 6 The **Change Security Key** screen is displayed. The cursor is at the **Security Key Identifier**. Enter an identifier for your Security Key.
- 7 Press <Tab> to enter a new passphrase.

 **NOTE:** You must provide the current passphrase in order to change the security key on the controller.

- 8 Press <Tab> and select **OK** to accept the settings and to exit the window. Select **Cancel** to exit if you do not want to change the security key on the controller.

 **NOTE:** If there is an existing configuration on the controller, it is updated with the new security key. If you had previously removed any secured disks, you still need to supply the old passphrase to import them.

Delete Key



NOTE: Delete Key is active if there is a security key present on the controller.



NOTE: Delete Key can only be performed when there are no secured virtual disks present.



NOTE: After the Delete Key operation, all unconfigured, secured SEDs shall be secure-erased.

Perform the following steps when deleting the security key on the controller:

- 1 During host system bootup, press <Ctrl><R> when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed. If there is more than one controller, the main menu screen is displayed.
- 2 Select a controller, and press <Enter>.
The **Virtual Disk Management** screen is displayed for the selected controller.
- 3 Use the arrow keys to highlight **Security Key Management**.
- 4 Press <F2> to display the actions you can perform.
- 5 Select the **Delete key** and press <Enter>.

Creating Secured Virtual Disks

To create a secured virtual disk, the controller must have a security key established first. See "Create Key" on page 122.



NOTE: Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and solid state drives (SSDs) within a virtual disk is not supported.

After the security key is established, perform the steps outlined in the "Creating Virtual Disks" on page 88 to create a virtual disk.

To secure the virtual disk, navigate to the **Secure VD** option at the bottom left area of the **Create New VD** screen.



NOTE: All virtual disks added to a secured Disk Group are secured.

Securing Pre-Existing Virtual Disks

If an unsecured virtual disk was created on a controller, you can secure the virtual disk as long as the following conditions are met:

- The controller has a security key present.
- All the physical disks in the virtual disk are SED.

Perform the following steps when securing a pre-existing virtual disk on a secured controller:

- 1 During the host system bootup, press <Ctrl><R> when the BIOS screen is displayed.

The **Virtual Disk Management** screen is displayed.

If there is more than one controller, the main menu screen is displayed.

- 2 Select a controller, and press <Enter>.

The **Virtual Disk Management** screen is displayed for the selected controller.

- 3 Use the arrow keys to highlight the **Disk Group** number.

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

- 5 Highlight the **Secure Disk Group** option and press <Enter>.



NOTE: If you select to secure a Disk Group, all virtual disks part of the Disk Group are secured.


Importing or Clearing Secured Foreign Configurations and Secure Disk Migration


Secured virtual disks created on a PERC H700 or H800 card can be migrated to another PERC H700 or H800 card. A virtual disk secured with a local security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them.


When importing secured virtual disk(s) created with a different security key, the secured foreign configuration(s) do not show in the **Foreign Configuration View** screen. Follow the steps below to import or clear a foreign secured virtual disk.



NOTE: If you are importing secured and unsecured virtual disks, you are prompted to resolve the secured foreign configuration first.

 **NOTE:** The PERC H700 or H800 card needs to have a security key present before being able to import a secured virtual disk.

 **NOTE:** Any unsecured virtual disks imported are still unsecured.

 **NOTE:** If you are importing a virtual disk originally secured with a local key (LKM), you are prompted for the passphrase used to secure that virtual disk.

Perform the following steps when importing a foreign secured virtual disk:

- 1 During the host system bootup, press <Ctrl><R> when the BIOS screen is displayed.

The **Virtual Disk Management** screen is displayed.


If there is more than one controller, the main menu screen is displayed.

- 2 Select a controller, and press <Enter>.

The **Virtual Disk Management** screen is displayed for the selected controller.

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

- 4 Select **Import** to import the foreign configuration or **Clear** to delete the foreign configuration. Press <Enter>.

 **NOTE:** To **Clear**, you need to **Instant Secure Erase** foreign configurations secured with a different security key.

- 5 If you select to **Import** the configuration, the **Secure Foreign Import** screen is displayed.

 **NOTE:** The key identifier for the passphrase used to secure the foreign secured virtual disks is displayed under the **Secured Drives** option.

- 6 Enter the passphrase that was used to secure the foreign configuration

- 7 Press <Tab> and select **OK** to finish importing the secured foreign configuration or select **Cancel** to exit this menu.

If you select **Cancel** for the secured foreign import, the disks remain inaccessible until imported or instant secure erased. See "Instant Secure Erase" on page 128.

Instant Secure Erase

Instant Secure Erase is the process of permanently erasing all data on an encryption-capable physical disk which is either unconfigured or foreign, and resetting the security attributes. You need to execute **Instant Secure Erase** on SEDs that are inaccessible (blocked) due to a lost or forgotten passphrase.

 **CAUTION: By executing Instant Secure Erase, the data on your encryption-capable physical disk is lost.**

Perform the following steps when instant secure erasing a physical disk:

- 1 Press <Ctrl><N> to access the **PD Mgmt** screen.
A list of physical disks is displayed. On the right menu, the physical disk properties are displayed including information about whether the physical disk is secured or not.
- 2 Press the down-arrow key to highlight a physical disk that is secured.
- 3 Press <F2> to display a menu of available actions.
- 4 The **Secure Erase** option is highlighted at the bottom of the menu.
- 5 Press <Enter> to Secure Erase the physical disk and select **YES**.

Troubleshooting Security Key Errors

Secured Foreign Import Errors

A foreign configuration is a RAID configuration that already exists on a replacement physical disk that you install in a system. A secured foreign configuration is a RAID configuration that was created under a different security key. There are three scenarios in which a secured foreign import fails:

- The passphrase authentication fails— A virtual disk secured with a security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them. Supply the correct passphrase to import the secured foreign configuration. If you have lost or forgotten the passphrase, the secured foreign disks remain locked (inaccessible) until the appropriate passphrase is entered or if they are instant secure erased.
- The secured virtual disk is in an offline state after supplying the correct passphrase— You must check to determine why the virtual disk failed and correct the problem. See "Troubleshooting" on page 131.

Failure to Select or Configure Non Self-Encrypting Disks (non-SED)

A virtual disk can be either secured or unsecured depending on how it was configured when created. In order to create a secured virtual disk, the controller must have a security key present and must be composed of SEDs only. In order to select/configure non-SED, you must create an unsecured virtual disk. You can create an unsecured virtual disk even if there is a security key present. Select the secure virtual disk option as **No** in the **Create New VD** menu. See "Creating Virtual Disks" on page 88 for steps on how to create an unsecured virtual disk.

Failure to Delete Security Key

A security key is used to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of data. If a security key is present, both secured and unsecured virtual disks may exist.

To delete the security key, you must have a previously established security key present on the controller and there can not be any configured secured disks. If there are configured secured disks, remove or delete them.

Failure to Instant Secure Erase Task on Physical Disks

Instant Secure Erase is the process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes. It is used in a couple of scenarios such as deleting a foreign configuration in the event of a forgotten or lost passphrase or unlocking a disk that had been previously locked.

Instant Secure Erase can be executed only on encryption-capable disks as long as the disks are not hot spares and are not configured (not part of a virtual disk). Ensure that the conditions are met and see "Instant Secure Erase" on page 128.

Troubleshooting

To get help with your Dell PowerEdge RAID Controller (PERC) H700 and H800 cards, you can contact your Dell Technical Service representative or access support.dell.com.

Post Error Messages

The controller BIOS read-only memory (ROM) provides Int 13h functionality (disk I/O) for the virtual disks connected to the controller. You can boot from or access the physical disks without a driver. Table 9-1 describes the error and warning messages for the BIOS.

Table 9-1. BIOS Errors and Warnings

Error Message	Probable Cause	Corrective Action
A discovery error has occurred, please power cycle the system and all the enclosures attached to this system	The message indicates that discovery did not complete within 120 seconds. The SAS cables for your system might be improperly connected.	Check the cable connections and restart the system.
There are X enclosures connected to connector Y, but only maximum of 4 enclosures can be connected to a single SAS connector. Please remove the extra enclosures then restart your system.	This message is displayed when the BIOS detects more than 4 enclosures connected to a single SAS connector.	Remove all additional enclosures and restart your system.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
Cache data was lost, but the controller has recovered. This could be due to the fact that your controller had protected cache after an unexpected power loss and your system was without power longer than the battery backup time. Press any key to continue or 'C' to load the configuration utility.	<p>This message is displayed under the following conditions:</p> <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 discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed.	<p>To resolve this problem, allow the battery to charge fully. If the issue persists, the battery or adapter DIMM might be faulty; contact Dell Technical Support.</p>
The following virtual disks have missing disks: (x). If you proceed (or load the configuration utility), these virtual disks will be marked OFFLINE and will be inaccessible. Please check your cables and ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility.	<p>The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.</p>	<p>Check the cable connections and restart the system.</p> <p>If there are no cable problems, press any key or <C> to continue.</p>

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
All of the disks from your previous configuration are gone. If this is an unexpected message, then please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility.	The message indicates that all configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.	Check the cable connections and restart the system. If there are no cable problems, press any key or <C> to continue.
The following virtual disks are missing: (x) If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility.	The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.	Check the cable connections and restart the system. If there are no cable problems, press any key or <C> to continue.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
The cache contains dirty data, but some virtual disks are missing or will go offline, so the cached data can not be written to disk. If this is an unexpected error, then please power off your system and check your cables to ensure all disks are present. If you continue, the data in cache will be permanently discarded. Press 'X' to acknowledge and permanently destroy the cached data.	The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.	Check the cable connections and restart the system. Use the <Ctrl><R> utility to import the virtual disk or discard the preserved cache. For the steps to manage preserved cache, see "Managing Preserved Cache" on page 96.
Entering the configuration utility in this state will result in drive configuration changes. Press 'Y' to continue loading the configuration utility or please power off your system and check your cables to ensure all disks are present and reboot.	The message is displayed after another BIOS warning indicating there are problems with previously configured disks and you have chosen to accept any changes and continue. The SAS cables for your system might be improperly connected.	Check the cable connections and restart the system. If there are no cable problems, press any key or <Y> to continue.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
BIOS Disabled. No Logical Drives Handled by BIOS	<p>This warning message is displayed after you disable the ROM option in the configuration utility.</p> <p>When the ROM option is disabled, the BIOS cannot boot to INT 13h and cannot provide the ability to boot from the virtual disk.</p> <p>Int 13h 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.</p>	Enable the ROM option.
Adapter at Baseport xxxx is not responding where xxxx is the baseport of the controller		Contact Dell Technical Support.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
There are offline or missing virtual drives with preserved cache. Please check the cables and ensure that all drives are present. Press any key to enter the configuration utility.	The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This preserved dirty cache is called pinned cache, and is preserved until you import the virtual disk, or discard the cache.	Use the <Ctrl><R> utility to import the virtual disk or discard the preserved cache. For the steps used to manage preserved cache, see "Managing Preserved Cache" on page 96.
x Virtual Disk(s) Offline where x is the number of virtual disks failed	This warning is displayed when the BIOS detects virtual disks in the offline state.	You must check to determine why the virtual disks failed and correct the issue. The BIOS does not take any action.
x Virtual Disk(s) Degraded where x is the number of virtual disks degraded	This message is displayed when the BIOS detects virtual disks in a degraded state.	Take corrective action(s) to make the virtual disks optimal. The BIOS does not take any action.
x Virtual Disk(s) Partially Degraded	This message is displayed when the BIOS detects a single disk failure in a RAID 6 or RAID 60 configuration.	You must check why the member disk is not present to correct the problem. The BIOS does not take any action.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
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 data in the controller cache that 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 discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed.• The battery may be under charged.	Allow the battery to charge fully to resolve this problem. If the issue persists, the battery or adapter DIMM might be faulty; contact Dell Technical Support.
Firmware is in Fault State		Contact Dell Technical Support.
Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility or 'F' to import foreign configuration(s) and continue.	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.	Press <F> at the prompt to import the configuration (if all member disks of the virtual disk are present) without loading the BIOS Configuration Utility. Or, press <C> to enter the BIOS Configuration Utility and either import or clear the foreign configuration.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
The foreign configuration message is present during POST but no foreign configurations are present in the foreign view page in CTRL+R. All virtual disks are in an optimal state.		<p>Ensure all your Physical Disks (PD) are present and all virtual disks are in optimal state. Clear the foreign configuration using <Ctrl><R> or Dell OpenManage Server Administrator Storage Management.</p> <p>⚠ CAUTION: The physical disk goes to Ready state when you clear the foreign configuration.</p> <p>If you insert a physical disk that was previously a member of a virtual disk in the system, and that disk's previous location has been taken by a replacement disk through a rebuild, you must manually remove the foreign configuration flag of the newly inserted disk.</p>
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 indicates that the controller and physical disks have different configurations.	You can use the BIOS Configuration Utility to clear the foreign configuration.

Table 9-1. BIOS Errors and Warnings (continued)



Error Message	Probable Cause	Corrective Action
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 restart the system.
Multibit ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press 'X' to continue or else power off the system, replace the controller and reboot.	This error is specific to the PERC H700 controller. Multi-bit ECC errors (MBE) occur in the memory and can corrupt cached data and discard it.	 CAUTION: MBE errors are serious, as they cause data corruption and data loss. In case of MBE errors, contact Dell Technical Support. NOTE: A similar message is displayed when multiple single-bit ECC errors are detected on the controller during bootup.
Multibit ECC errors were detected on the RAID controller. The DIMM on the controller needs replacement. If you continue, data corruption can occur. Press 'X' to continue or else power off the system, replace the DIMM module, and reboot. If you have replaced the DIMM please press 'X' to continue.	Multi-bit ECC errors (MBE) occur in the memory and can corrupt cached data and discard it.	 CAUTION: MBE errors are serious, as they cause data corruption and loss. In case of MBE errors, contact Dell Technical Support.

Table 9-1. BIOS Errors and Warnings (continued)

Error Message	Probable Cause	Corrective Action
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.	<p>The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible.</p> <p>The SAS cables for your system might be improperly connected.</p>	<p>Check the cable connections and restart the system.</p> <p>If there are no cable problems, press any key or <C> to continue.</p>
The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging then contact technical support for additional assistance.	<p>NOTE: This message may appear for controller without battery, depending on the virtual disks' policies.</p>	<ul style="list-style-type: none">• The controller battery is missing or currently inaccessible. Contact Dell Technical Support if the issue persists after 30 minutes.• The controller battery is completely discharged and needs to be charged for it to become active. You must first charge the battery, then wait for a few minutes for the system to detect it.

Degraded State of Virtual Disks

A redundant virtual disk is in a degraded state when one or more physical disks have failed or are inaccessible. For example, if a RAID 1 virtual disk consists of two physical disks and one of them fails or become inaccessible, the virtual disk become degraded.

To recover a virtual disk from a degraded state, you must replace the failed physical disk and rebuild it. Once the rebuilding process is complete, the virtual disk state changes from **degraded** to **optimal**. For information on rebuilding the disk, see "Performing a Manual Rebuild of an Individual Physical Disk" on page 111.

Memory Errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from the memory errors. Single-bit memory errors can be handled by the controller and do not disrupt normal operation. A notification is 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 a multi-bit error occurs while accessing data in the cache when the controller is started with dirty cache, the controller discards the cache contents. The controller generates a warning message to the system console to indicate that the cache was discarded and generates an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the controller stops.
- The controller logs an event to the controller's internal event log and a message during POST is displayed indicating a multi-bit error has occurred.



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

Preserved Cache State

The controller preserves the dirty cache from a virtual disk if the virtual disk becomes offline or is deleted because of missing physical disks. This preserved dirty cache is called pinned cache and is preserved until you import the virtual disk or discard the cache.

Use the BIOS Configuration Utility (<Ctrl><R>) utility to select whether to import the virtual disk or discard the preserved cache. In the **VD Mgmt** menu, select **Manage Preserved Cache** and follow the steps on the screen.

General Issues

Table 9-2 describes general problems you might encounter, with suggested solutions.

Table 9-2. General Problems

Issue	Corrective Action
The device is displayed in Device Manager but has a yellow bang (exclamation mark).	Reinstall the driver. For more information on reinstalling drivers, see "Driver Installation" on page 69.
The device does not appear in the Device Manager .	Turn off the system and reseal the controller.
No Hard Drives Found message is displayed during a media-based installation of Microsoft Windows Server 2003 because of the following causes:	The corresponding solutions are:
<ul style="list-style-type: none">• The driver is not native in the operating system.• The virtual disks are not configured properly.• The controller BIOS is disabled.	<ul style="list-style-type: none">• Press <F6> to install the RAID device driver during installation.• Enter the BIOS Configuration Utility to configure the virtual disks. For procedures to configure the virtual disks, see the section "Configuring and Managing RAID" on page 81.• Enter the BIOS Configuration Utility to enable the BIOS. For information on configuring virtual disks, see "Installing and Configuring Hardware" on page 41.

Physical Disk Related Issues

Table 9-3 describes physical disk-related problems you might encounter, with suggested solutions.

Table 9-3. Physical Disk Issues

Issue	Corrective Action
One of the physical disks in the disk array is in the failed state.	Update the PERC H700 and H800 cards to the latest firmware available on support.dell.com .
Cannot rebuild a fault tolerant virtual disk. NOTE: For more information, see the alert log for virtual disks.	The replacement disk is too small or not compatible with the virtual disk. Replace the failed disk with a compatible good physical disk with equal or greater capacity.
Fatal error(s) or data corruption(s) are reported when accessing virtual disks.	Contact Dell Technical Support.
One or more physical disks is displayed as Blocked and can not be configured.	Update the PERC H700 and H800 cards to the latest firmware available on support.dell.com .

Physical Disk Failures and Rebuild Issues

Table 9-4 describes issues related to physical disk failures and rebuilds.

Table 9-4. Physical Disk Failure and Rebuild Issues


Issue	Corrective Action
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.</p> <p>You can recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk:</p>
	<p> CAUTION: Follow the safety precautions to prevent electrostatic discharge.</p>
	<ol style="list-style-type: none">1 Turn off the system, check cable connections, and reseal physical disks.2 Ensure that all the disks are present in the enclosure.3 Turn on the system and enter the <Ctrl><R> utility and import the foreign configuration. Press <F> at the prompt to import the configuration, or press <C> to enter the BIOS configuration utility and either import or clear the foreign configuration.
	<p>If the virtual disk is redundant and transitioned to Degraded state before going Offline, a rebuild operation starts automatically after the configuration is imported.</p> <p>If the virtual disk has gone directly to the Offline state due to a cable pull or power loss situation, the virtual disk is imported in its Optimal state without a rebuild occurring.</p>
	<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>For information on rebuilding a single physical disk, see "Performing a Manual Rebuild of an Individual Physical Disk" on page 111.</p>

Table 9-4. Physical Disk Failure and Rebuild Issues (continued)

Issue	Corrective Action
Rebuilding a physical disk after one of them is in a failed state.	If you have configured hot spares, the PERC H700 or PERC H800 card automatically tries to use one of them to rebuild a physical disk that is in a failed state. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed physical disks are available. You must insert a physical disk with enough storage in 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. For information on rebuilding a single physical disk, see "Performing a Manual Rebuild of an Individual Physical Disk" on page 111.
A virtual disk fails during rebuild while using a global hot spare.	The global hot spare goes back to Hotspare state and the virtual disk goes to Failed state.
A virtual disk fails during rebuild while using a dedicated hot spare.	The dedicated hot spare goes to Ready state and the virtual disk goes to Failed state.
A physical disk fails 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 I/O operation for every five host I/O operations.
You cannot add a second virtual disk to a disk group while the virtual disk in that disk group is undergoing a rebuild	The firmware does not allow you to create a virtual disk using the free space available in a disk group if a physical disk in a virtual disk group is undergoing a rebuild operation.

SMART Errors

Table 9-5 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 at support.dell.com/manuals.

Table 9-5. SMART Errors

Issue	Corrective Action
A SMART error is detected on a physical disk in a redundant virtual disk.	<p>Perform the following steps:</p> <ol style="list-style-type: none">1 Back up your data.2 Force the physical disk offline. <p>NOTE: If a hot spare is present, the rebuild starts with the hot spare after the disk is forced offline.</p> <ol style="list-style-type: none">3 Replace it with a new physical disk of equal or higher capacity.4 Perform the Replace Member operation. <p>The Replace Member operation allows you to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. For more information about the Replace Member feature, see "Using Replace Member and Reversible Hot Spares" on page 35.</p>
A SMART error is detected on a physical disk in a non-redundant virtual disk.	<p>Perform the following steps:</p> <ol style="list-style-type: none">1 Back up your data.2 Use Replace Member or set up a global hot spare to replace the disk automatically. <p>For more information about the Replace Member feature, see "Using Replace Member and Reversible Hot Spares" on page 35.</p> <ol style="list-style-type: none">3 Replace the affected physical disk with a new physical disk of equal or higher capacity.4 Restore from the backup.

Replace Member Errors

Table 9-6 describes issues related to the **Replace Member** feature.



NOTE: For more information about the **Replace Member** features, see "Using Replace Member and Revertible Hot Spares" on page 35.

Table 9-6. Replace Member Operation Errors

Issue	Corrective Action
The source disk fails during the Replace Member operation.	If the source data is available from other disks in the virtual disk, the rebuild begins automatically on the target disk, using the data from the other disks.
Target disk fails	If the target disk fails, the Replace Member operation aborts.
Other disks fail	If the target disk fails and the Replace Member operation aborts but the source data is still available, then the Replace Member operation continues as Replace Member .

Linux Operating System Errors

Table 9-7 describes issues related to the Linux operating system.

Table 9-7. Linux Operating System Errors

Error Message	Corrective Action
<pre><Date:Time> <HostName> kernel: sdb: asking for cache data failed</pre>	<p>This error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings.</p>
<pre><Date:Time> <HostName> kernel: sdb: assuming drive cache: write through</pre>	<p>The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. 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 "Physical Disk Management" on page 108 for more information about Write-Through cache.</p>
	<p>Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC H700 and PERC H800 SAS RAID system remain unchanged.</p>

Table 9-7. Linux Operating System Errors (continued)

Error Message	Corrective Action
Driver does not auto-build into new kernel after customer updates.	<p>This error is a generic issue for Dynamic Kernel Module Support (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"><li data-bbox="526 399 1006 430">1 Install a DKMS-enabled driver package.<li data-bbox="526 438 1006 494">2 Run <code>up2date</code> or a similar tool to upgrade the kernel to the latest version.<li data-bbox="526 502 1006 534">3 Reboot to the new kernel. <p>The driver running in the new kernel is the native driver of the new kernel. The driver package you installed previously 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"><li data-bbox="526 734 1006 861">1 Type: <pre>dkms build -m <module_name> -v <module version> -k <kernel version></pre><li data-bbox="526 869 1006 997">2 Type: <pre>dkms install -m <module_name> - v <module version> -k <kernel version></pre><li data-bbox="526 1005 1006 1235">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 9-7. Linux Operating System Errors (continued)

Error Message	Corrective Action
smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device	This is a known issue. An unsupported command is entered through the user application. User applications attempt to direct Command Descriptor Blocks to RAID volumes. The error message does not effect the feature functionality.
smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf	The Mode Sense/Select command is supported by firmware on the controller. However, the Linux kernel daemon issues the command to the virtual disk instead of to the driver IOCTL node. This action is not supported.

Disk Carrier LED Indicators

The LED on the physical disk carrier indicates the state of each physical disk. Each disk carrier in your enclosure has two LEDs: an activity LED (green) and a status LED (bicolor, green/amber) as shown in Figure 9-1. The activity LED is active whenever a disk is being accessed while the status LED indicates when a disk is being spun up, is rebuilding, or is in a fault state.

Figure 9-1. Disk Carrier LED Indicators



1 activity LED

2 status LED



Regulatory Notices

For additional regulatory information, please go to the Regulatory Compliance Homepage on dell.com at the following location: dell.com/regulatory_compliance.

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害和/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。

“0”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。



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 products addressed by this document:

Dell B.V. Taiwan Branch

20/F, No. 218, Sec. 2, Tung Hwa S. Road,
Taipei, Taiwan

Glossary

A

Adapter Card

An adapter card enables the system to access peripheral devices by converting the protocol of one bus or interface to another. For example, a RAID controller is a type of adapter card that provides RAID functions. Adapter cards may reside on the system board. It may also be in the form of an add-in card.

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 occur 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.

Array

A grouping of physical disks that combines the storage space on the physical disks into a single segment of contiguous storage space. The RAID controller can group physical disks on one or more channels into an array. A hot spare disk does not participate in an array.

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

Baseport is the 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

Basic Input/Output System. Your system'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, also known as <Ctrl><R>, configures and maintains RAID disk groups and virtual disks, and manages the RAID system. The operation of this utility is independent because the utility resides in the controller BIOS.

C

Cache

A 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. For more information, see "Write-Back" on page 173.

Change Key

The process of generating a key for an encryption-capable or a security-capable component. All current data on the media is accessible using the newly generated key.

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, 6, 50, and 60 arrays, consistency check verifies correct parity data for each stripe. For RAID 1 and RAID 10 arrays, this operation verifies correct mirror data for each stripe.

Controller

See "Adapter Card" on page 157

Controller-Bound Security Configuration

An encryption configuration whereby the controller encrypts the security key using an internal algorithm before storing it internally in the NVData. The controller always understands how to decrypt the key and you need not provide a passphrase on boot.

Current Encryption State

The state of an encryption-capable component. States are encrypted and not encrypted.

Current Security State

The state of a security-capable component. States are secured and not secured.

Decrypt

The process of reversing the obfuscation of data with the full knowledge of the algorithm and the key used to encrypt it.

Default Encryption State

The encryption state to which an encryption enabled component reverts at power-on (or after an internal reset such as a firmware upgrade). Encryption states are of two types: encrypted and not encrypted.

Default Security State

The security state to which a security enabled component is reverted at power-on (or after an internal reset such as a firmware upgrade). Security states are of two types: secured and not secured.

D**DDR SDRAM**

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 collection of disks from one or more disk subsystems combined using a configuration utility. The utility controls the disks and presents them to the array operating environment as one or more logical drives.

Disk Group

A logical grouping of disks attached to a RAID controller on which one or more virtual disks can be created.

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 data is 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

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

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

E**ECC Errors**

Error Correcting Code. It refers to errors detected during memory transactions. Single-bit ECC errors can be handled by the firmware and do not disrupt normal operation. ECC double-bit or multi-bit errors are more serious, as they cannot be corrected and may result in the controller becoming inoperable.

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.

Encrypt

The act of obfuscating data on a media component through an algorithm that relies on an encryption key.

Encrypted

State of an encryption-enabled media component indicating that its data has been obfuscated using an algorithm that relies on an encryption key.

Encryption Capable

Indicates that a component can be encrypted.

Encryption Enabled/Disabled

Indicates whether an encryption capable media component is encrypted.

Encryption Key

The key used to obfuscate the data in an encryption enabled media component.

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 disk failure per disk group without compromising data integrity and processing capability. The PERC H700 or PERC H800 cards provide this support through redundant virtual disks in RAID levels 1, 5, 6, 10, 50, and 60.

Firmware

Software stored in read-only memory (ROM) or Programmable ROM (PROM).

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.

G

GB

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. Servers and workstations can 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.

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 passes a consistency check.

Instant Secure Erase

The process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes.

K**Key Identifier**

The user supplied clear text label used to associate the correct security key with the controller. Examples: Encryption Key Identifier, Security Key Identifier.

L**Load-Balancing**

Load Balancing, a feature that is enabled with a redundant path configuration on the PERC H800 card, ensures that a balanced number of disks are on each port of the enclosure. The load-balancing architecture can also provide additional performance increases in some workloads where the required bandwidth can exceed that of what a single SAS link can provide.

M

MB

Megabyte(s). The term *megabyte* means 1,048,576 bytes (2^{20} bytes); however, when referring to disk 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.

N

No Read-Ahead

No read ahead is a cache read policy. If you select No 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 is needed soon. No 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 is a non redundant virtual disk and 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.

NVRAM

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.

Non-Volatile Cache

A cache module with flash-based storage to preserve cache data indefinitely. This replaces the need for a battery backup unit (BBU) to supply power to preserve cache data.

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.

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.

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 disk failure and damage data integrity.

Passphrase

The user supplied string that the controller uses to create the security key

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 SATA device.

Physical Disk

A non-volatile, randomly addressable device for storing data. Physical disks are rewritable and can also referred to as hard drives and solid state drives (SSDs).

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**

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, 6, 10, 50, and 60) 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 you.

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 Dell OpenManage storage management application 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 after you have installed the operating system.

SAS RAID Storage Manager configures, monitors, and maintains the PERC H700 card, battery backup units, and other devices running on a system.

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 is required 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, 6, 10, 50, or 60) 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. 100% rebuild rate does not mean ALL CPU resources are dedicated to the rebuild without processing IOs.

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 Path

The PERC H800 firmware provides support for detecting and establishing redundant paths from the RAID controller to the SAS devices in the enclosure. With redundant paths, if one path fails, another path can be used to maintain communication between the controller and the enclosure.

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.

Replace Member

The procedure used to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. The **Replace Member** operation is often used to create or restore a specific physical configuration for an array (for example, a specific arrangement of array members on the device I/O buses).

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.

Reversible Hot Spare

In the case of a Reversible Hot Spare, when you use the **Replace Member** procedure, after data is copied from a hot spare to a new physical disk, the hot spare reverts from a rebuild disk to its original hot spare status.

RPM

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

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.

SATA

Serial Advanced Technology Attachment. A physical storage interface standard, is a serial link that provides point-to-point connections between devices.

Secure

The act of creating a security key to control access to encrypted data.

Secured

State of a security-enabled media component indicating that its data is only accessible using a security key.

Related term: Not Secured

Secure Disk Group

Indicates that a disk group comprised of security-capable physical disks is secured.

Secure Migration

The process of moving a set of secured disks to a different controller.

Security Enabled/Disabled

Indicates whether a security capable component is secured.

Security Key

A key to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of the data.

Self Encrypting Disk (SED)

Hardware-based encryption incorporated in a security-capable physical disk.

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 is sent if the number of single-bit errors exceeds a threshold value.

SMART

The self-monitoring analysis and reporting technology (SMART) feature monitors the internal performance of all motors, heads, and disk electronics to detect predictable disk failures.

SMP

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.

Spanning

The method by which nested RAID levels (such as RAID 10, 50, and 60) are constructed from multiple sets of basic, or single RAID levels. For example, a RAID 10 is made up of multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. The same concept holds true for RAID 50 and 60 where multiple sets of RAID 5 or RAID 6 can be combined together with RAID levels.

Spare

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

SSD

Solid-state Disk. SSDs are storage devices that use solid-state memory to store data as opposed to traditional rotational hard drives. SSDs can be made to use either the SAS or SATA protocol.

SSP

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.

STP

Serial Tunneling Protocol, STP, enables communication with a 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.

U

User-Bound Security Configuration

A configuration in which the controller encrypts the security key with the passphrase and then stores it in the NVData. The controller cannot decrypt the security key without the passphrase. The controller prompts the user for the passphrase on every boot.

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. The policies include the amount of dirty or 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.

Index

B

- Background Initialization, 25
- background initialization, 25, 157
 - stopping, 111
- baseport, 157
- battery
 - management, 28
 - removing from PERC 5/E Adapter, 64
- BIOS, 107, 158
- BIOS Configuration
 - Utility, 106-107, 158
 - controller management, 107
 - Foreign View menu, 108
 - menu navigation controls, 83
 - menu options, 99
 - menus, 99
 - physical disk management, 105
 - starting, 82
 - virtual disk management, 100

C

- cache, 158
- compatibility
 - with existing RAID controllers, 24
- Consistency Check, 26

- consistency check, 91, 104, 159
- controller, 159, 169
- controller management, 107

D

- disk groups
 - deleting, 98
- disk migration, 27
- disk mirroring, 19
- Disk roaming, 26
- disk roaming, 26
- disk striping, 18
- display/update parameters, 104
- distributed parity, 161
- driver diskette, 69
- driver installation, 69
- drivers
 - installation, 69
 - Microsoft operating system installation, 71

E

- electrostatic discharge. *See* ESD
- ESD, 12

F

- fault tolerance, 32
 - features, 32
- foreign configuration, 108
- Foreign Configuration View, 108
- full initialization, 25

H

- hot swap, 164
- hot swapping, 34

I

- initialization, 164
- interface specifications, 24

L

- LED
 - operation, 26

M

- manual rebuild, 111

O

- operating system support, 16
- operating systems, 16

P

- parity, 19, 166
 - distributed, 161
- Patrol Read, 38
- PCI
 - architecture, 16
- PERC
 - overview, 15
- PERC 6
 - controller descriptions, 15
- PERC H700, H800
 - Card Descriptions, 15
 - supported operating systems, 16
- physical disk
 - actions, 106
- physical disks
 - actions, 106
 - management, 105
- post error messages, 131

R

- RAID, 164
 - configuration, 81
 - configuration and management, 81
 - configuration functions, 85
 - configuration information, 31
 - definition, 17, 167
 - description, 17
 - level migration, 168
 - levels, 17, 104, 168

- management, 81
- summary of levels, 17

RAID level, 86

RAID levels, 168

read policy, 87

rebuild, 106

- manual, 111

Red Hat Enterprise Linux

- creating a driver diskette, 74

- installing with the driver update diskette, 76

replacement disk, 169

S

safety instructions

- for preventing ESD, 12

SAS, 170

- controller descriptions, 15

- overview, 15

SATA, 170

SMART technology, 24

spare, 172

stripe element size, 86

T

troubleshooting, 131

- general problems, 142

- physical disk issues, 143

- post error messages, 131

SMART error, 146-147

V

virtual disks

- deleting, 98

- management, 100

- menu options, 104, 107

- parameters, 86

- setting up, 107-108

W

write policy, 87

