

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

[Safety Instructions](#)

[Overview](#)

[Introduction to RAID](#)

[Features](#)

[RAID Configuration and Management](#)

[Driver Installation](#)

[Troubleshooting](#)

[Appendix A: Regulatory Notice](#)

[Glossary](#)

Information in this document is subject to change without notice.
© 2003-2005 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, and *Dell OpenManage* are trademarks of Dell Inc. *MegaRAID* is a registered trademark of LSI Logic Corporation. *Microsoft*, *Windows NT*, *MS-DOS*, and *Windows* are registered trademarks of Microsoft Corporation. *Intel* is a registered trademark of Intel Corporation. *Novell* and *NetWare* are registered trademarks of Novell, Inc. *RedHat* is a registered trademark of RedHat, Inc.

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

Model PERC 4/Di/Si, PERC 4e/Di/Si

Release: April 2005

Part Number: GC687 Rev.A07

[Back to Contents Page](#)

Appendix A: Regulatory Notice

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [FCC Notices \(U.S. Only\)](#)
 - [A Notice About Shielded Cables](#)
 - [Class B](#)
 - [Canadian Compliance \(Industry Canada\)](#)
 - [MIC Notice \(Republic of Korea Only\)](#)
 - [VCCI Class B Statement](#)
-

FCC Notices (U.S. Only)

Most Dell systems are classified by the Federal Communications Commission (FCC) as Class B digital devices. However, the inclusion of certain options changes the rating of some configurations to Class A. To determine which classification applies to your system, examine all FCC registration labels located on the back panel of your system, on card-mounting brackets, and on the controllers -themselves. If any one of the labels carries a Class A rating, your entire system is considered to be a Class A digital device. If all labels carry either the Class B rating or the FCC logo (FCC), your system is considered to be a Class B digital device.

Once you have determined your system's FCC classification, read the appropriate FCC notice. Note that FCC regulations provide that changes or modifications not expressly approved by Dell Inc. could void your authority to operate this equipment.

A Notice About Shielded Cables

Use only shielded cables for connecting peripherals to any Dell device to reduce the possibility of interference with radio and television reception. Using shielded cables ensures that you maintain the appropriate FCC radio frequency emissions compliance (for a Class A device) or FCC certification (for a Class B device) of this product. For parallel printers, a cable is available from Dell Inc.

Class B

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

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

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

If necessary, consult a representative of Dell Inc. or an experienced radio/television technician for additional suggestions. You may find the following booklet helpful: FCC Interference Handbook, 1986, available from the U.S. Government Printing Office, Washington, DC 20402, Stock No. 004-000-00450-7. This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

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

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

- 1 Product name: Dell PowerEdge Expandable RAID Controller 4 Controller
- 1 Company name: Dell Inc.

Regulatory Department

One Dell Way

Round Rock, Texas 78682 USA

512-338-4400

Canadian Compliance (Industry Canada)

Canadian Regulatory Information (Canada Only)

This digital apparatus does not exceed the Class B limits for radio noise emissions from digital apparatus set out in the Radio Interference Regulations of the Canadian Department of Communications. Note that the Canadian Department of Communications (DOC) regulations provide, that changes or modifications not expressly approved by Intel could void your authority to operate the equipment. This Class B digital apparatus meets all the requirements of the Canadian Interference -Causing Equipment Regulations.

Cet appareil numerique de la classe B respecte toutes les exigences du Reglement sur la material brouilleur du Canada.

MIC Notice (Republic of Korea Only)

B Class Device

기종별	사용자 안내문
B급 기기 (가정용 정보통신기기)	이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든 지역에서 사용할 수 있습니다.

Please note that this device has been approved for non-business purposes and may be used in any environment, including residential areas.



기기명칭 PCI-E Host Adapter
 기기등록번호 01097
 증명 또는 상호 I.S.I.I. CORP. CORPORATION
 제조자 및 제조국가 SCS SYSTEMS (THAILAND) ,TD./THAI
 인증년월일 2004년(Year) 03월(Month) 19일(Date)



기 기 크 명 칭	PCI Host Adapter
기 기 의 도 과 명	SERIES 618 and 620
성 령 투 주 상 호	LSI LOGIC CORPORATION
제 조 자 및 제 조 국 가	LSI SYSTEMS (THAILAND) LTD./미국
인 승 연 도 월 일	2004년(Year) 09월(Month) 15일(Date)

VCCI Class B Statement

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

取扱説明書に従って正しい取り扱いをして下さい。

[Back to Contents Page](#)

[Back to Contents Page](#)

Overview

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [Features](#)
- [RAID and SCSI Modes](#)
- [Changing the Mode on the Embedded RAID Controller from RAID/RAID to RAID/SCSI Mode or from RAID/SCSI to RAID/RAID Mode](#)

The Dell™ PowerEdge™ Expandable RAID Controller (PERC) 4/Di/Si and 4e/Di/Si are embedded subsystems on the motherboard that offer RAID control capabilities. The RAID controller supports all low-voltage differential (LVD) SCSI devices on Ultra320 and Wide SCSI channels with data transfer rates up to 320 MB/sec. PERC 4/Si and 4e/Si support a single channel, while PERC 4/Di and 4e/Di support two channels.

The RAID controller provides reliability, high performance, and fault-tolerant disk subsystem management. It is an ideal RAID solution with Dell PowerEdge systems. This RAID controller offers a cost-effective way to implement RAID in a server and provides reliability, high performance, and fault-tolerant disk subsystem management.

Features

The RAID controller features include:

- 1 Wide Ultra320 LVD SCSI performance of up to 320 MB/s
- 1 Support for 256 MB (DDR2) memory
- 1 64-bit/66 MHz peripheral component interconnect (PCI) host interface for PERC 4/Di/Si
- 1 PCI Express x8 host interface for PERC 4e/Di/Si
- 1 RAID levels 0 (striping), 1 (mirroring), 5 (distributed parity), 10 (combination of striping and mirroring), and 50 (combination of striping and distributed parity)
- 1 Advanced array configuration and management utilities
- 1 Ability to boot from any array
- 1 One electrical bus: a LVD bus.

 **NOTE:** The PERC 4/Di/Si and 4e/Di/Si RAID controllers support hard disk drives only; they do not support CD-ROMs, tape drives, tape libraries, or scanners.

Hardware Architecture

PERC 4/Di/Si supports a (PCI) host interface, while PERC 4e/Di/Si supports PCI Express x8 host interface. PCI-Express are high-performance I/O bus architectures designed to increase data transfers without slowing down the CPU. PCI-Express goes beyond the PCI specification in that it is intended as a unifying I/O architecture for various systems: desktops, workstations, mobile, server, communications, and embedded devices.

Maximum Cable Length for 320M SCSI

The maximum length of cable that you can use for LVD 320M SCSI is 12 meters (39' 4"), with a maximum number of 15 devices.

Operating System Support

The RAID controller supports the following operating systems:

- 1 Microsoft® Windows® 2000: Server, AS
- 1 Windows Server 2003: Standard Edition, Enterprise Edition, Small Business Edition
- 1 Novell® NetWare®

 **NOTE:** See [Driver Installation](#) for the latest operating system versions and driver installation procedures for the operating systems.

RAID and SCSI Modes

RAID mode allows the channel on the controller to support RAID capabilities, while SCSI mode allows the channel to operate as a SCSI channel. Devices attached to the SCSI channel are not controlled by the RAID firmware and function as if attached to a regular SCSI controller. Check your system documentation to determine supported modes of operation.

You can use system setup to select RAID or SCSI mode. During bootup, press <F2> to access system setup. The PERC 4/Si and 4e/Si RAID controllers work with cache memory and a card key to support one channel, which can be in either SCSI or RAID mode.

The PERC 4/Di and 4e/Di RAID controllers work with cache memory and a card key to provide two SCSI channels to support configurations that span internal channels and external enclosure channels. You can use available physical drives to make a logical drive (volume). The drives can be on different channels, internal or external.

 **NOTE:** The maximum number of drives you can use depends on your system configuration.

For PERC 4/Di and 4e/Di, [Table 1-1](#) displays the possible combinations of SCSI and RAID modes for channels 0 and 1 on the controller.

Table 1-1. SCSI and RAID Modes for the PERC 4/Di and 4e/Di RAID Controller

Mode	Channel 0	Channel 1
RAID	RAID	RAID
RAID/SCSI (if supported by your platform)	RAID	SCSI
SCSI	SCSI	SCSI

 **NOTE:** You cannot set Channel 0 as SCSI while Channel 1 is set as RAID.

Use the mixed mode (RAID on channel 0, SCSI on channel 1, known as *RAID/SCSI* mode), where available, with a RAID channel for hard drives and a legacy SCSI channel for removable devices or pre-existing hard drives, where available. Not all systems support RAID/SCSI mode.

If both channels are in SCSI mode, you can change channel 0 to RAID to create a RAID/SCSI mode. Dell recommends that you leave the SCSI channel that contains the operating system in SCSI mode. However, you cannot leave channel 0 as SCSI and change channel 1 to RAID, because SCSI/RAID mode is not allowed.

 **NOTICE:** You will lose data if the configuration is changed from SCSI to RAID, RAID to SCSI, or RAID/RAID to RAID/SCSI.

 **NOTE:** SCSI/SCSI is not a RAID configuration and is available only if you disable RAID by selecting SCSI mode in the system BIOS. (Access the system BIOS by pressing <F2> during bootup). See your systems *User's Guide* to learn how to select SCSI and RAID modes in the system BIOS.

Drive Size in RAID and SCSI Modes

The capacity of a hard drive is reported differently when the hard drive is on the SCSI channel of a PERC 4/Di or 4e/Di controller in RAID/SCSI mode and SCSI/SCSI mode.

The size reported by firmware while in SCSI mode is the actual size in megabytes. For example, a hard drive size of 34734 MB is 36,422,000,000 bytes divided by 1048576 (1024 * 1024, the actual number of bytes in 1 MB), which is off by 2 MB.

In RAID mode, the coerced size is rounded down to the nearest 128 MB boundary, then rounded to the nearest 10 MB boundary. Drives in the same capacity class, such as 36 GB, but from different vendors usually do not have the exact same physical size. With drive coercion, the firmware forces all the drives in the same capacity class to the same size. This way, you can replace a larger drive in a class with a smaller drive in the same class.

Changing the Mode on the Embedded RAID Controller from RAID/RAID to RAID/SCSI Mode or from RAID/SCSI to RAID/RAID Mode

The embedded RAID controller on the system supports two modes of operations: RAID/RAID and RAID/SCSI. The RAID/RAID mode allows the system to use both SCSI channels for RAID only operation. The RAID/SCSI mode allows the system to use RAID for the internal SCSI disk drives and reserves one SCSI channel to allow the connection of internal tape or external SCSI devices. Before changing from RAID/RAID to RAID/SCSI (or from RAID/SCSI to RAID/RAID), you must manually clear the RAID configuration to avoid configuration problems.

 **NOTE:** Dell does not support changing from RAID/RAID to RAID/SCSI or from RAID/SCSI to RAID/RAID with RAID virtual disks already created. If you change mode without clearing the RAID configuration, then you may experience unexpected behavior on the system or system instability.

The following procedures are required when changing the embedded RAID controller from RAID/RAID to RAID/SCSI mode or from RAID/SCSI to RAID/RAID mode:

 **NOTICE:** These steps will delete all data on the hard drives. Back up all your data before proceeding.

Clearing the controller configuration:

1. Reboot the system.
2. When the RAID controller initialization displays, press <Ctrl> <M> to enter the RAID controller configuration utility.

If your system has add-on RAID controllers in addition to the embedded RAID controller, proceed to [step 3](#). If your system has only the embedded RAID controller, then proceed to [step 5](#).

3. Select **Select Adapter**.
4. Select the embedded RAID controller and press <Enter>.
5. Select **Configure**.
6. Select **Clear Configuration**.

 **NOTICE:** All your data will be lost after you perform this step. Do not perform this step until you have backed up all your data.

7. Select **Yes** to confirm.
8. Press any key to return to the menu.
9. Press <ESC> twice to exit the menu.
10. When prompted to exit, select **Yes** to exit the menu.
11. Reboot the system.

Changing the RAID Mode

1. Press <F2> to enter the system BIOS configuration.
2. Select **Integrated devices** and press <Enter> to enter the **Integrated Devices** menu.
3. Move your selection to **Channel B** under **Embedded RAID controller**.
 - a. To change from RAID/RAID to RAID/SCSI, change this value from RAID to SCSI.
 - b. To change from RAID/SCSI to RAID/RAID, change this value from SCSI to RAID
4. Press <ESC> to exit the Integrated Devices Menu.
5. Press <ESC> again to exit the BIOS and reboot the system.

During the system boot, the following warning message displays to confirm the mode change:

```
Warning: Detected mode change from RAID to SCSI (or from SCSI to RAID) on channel B of the embedded RAID subsystem.
```

```
Data loss will occur!
```

6. Press <Y> to confirm this change.

7. Press <Y> again to verify the change.

Recreating your RAID configuration:

1. When the RAID controller initialization displays, press <Ctrl><M> to enter the RAID controller configuration utility.
2. Create the RAID volumes required for your desired configuration.

 **NOTE:** Refer to [RAID Configuration and Management](#) for more information on how to create RAID volumes using the RAID controller configuration utility.

[Back to Contents Page](#)

[Back to Contents Page](#)

Introduction to RAID

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [Components and Features](#)
- [RAID Levels](#)
- [RAID Configuration Strategies](#)
- [RAID Availability](#)
- [Configuration Planning](#)

RAID is an array of multiple independent hard disk drives that provides high performance and fault tolerance. The RAID array appears to the host computer as a single storage unit or as multiple logical units. Data throughput improves because several disks can be accessed simultaneously. RAID systems also improve data storage availability and fault tolerance. Data loss caused by a hard drive failure can be recovered by reconstructing missing data from the remaining data or parity drives.

RAID Description

RAID (Redundant Array of Independent Disks) is an array, or group, of multiple independent hard drives that provide high performance and fault tolerance. A RAID disk subsystem improves I/O (input/output) performance and reliability. The RAID array appears to the host computer as a single storage unit or as multiple logical units. I/O is expedited because several disks can be accessed simultaneously.

RAID Benefits

RAID systems improve data storage reliability and fault tolerance compared to single-drive storage systems. Data loss resulting from a hard drive failure can be prevented by reconstructing missing data from the remaining hard drives. RAID has gained popularity because it improves I/O performance and increases storage subsystem reliability.

RAID Functions

Logical drives, also known as virtual disks, are arrays or spanned arrays that are available to the operating system. The storage space in a logical drive is spread across all the physical drives in the array.

 **NOTE:** The maximum logical drive size for all supported RAID levels (0, 1, 5, 10, and 50) is 2 TB. You can create multiple logical drives on the same physical disks.

Your SCSI hard drives must be organized into logical drives in an array and must be able to support the RAID level that you select. Below are some common RAID functions:

- 1 Creating hot spare drives.
- 1 Configuring physical arrays and logical drives.
- 1 Initializing one or more logical drives.
- 1 Accessing controllers, logical drives, and physical drives individually.
- 1 Rebuilding failed hard drives.
- 1 Verifying that the redundancy data in logical drives using RAID level 1, 5, 10, or 50 is correct.
- 1 Reconstructing logical drives after changing RAID levels or adding a hard drive to an array.
- 1 Selecting a host controller to work on.

Components and Features

RAID levels describe a system for ensuring the availability and redundancy of data stored on large disk subsystems. PERC 4/Di/Si and 4e/Di/Si support RAID levels 0, 1, 5, 10 (1+0), and 50 (5+0). See [RAID Levels](#) for detailed information about RAID levels.

Physical Array

A physical array is a group of physical disk drives. The physical disk drives are managed in partitions known as logical drives.

Logical Drive

A logical drive is a partition in a physical array of disks that is made up of contiguous data segments on the physical disks. A logical drive can consist of an entire physical array, more than one entire physical array, a part of an array, parts of more than one array, or a combination of any two of these conditions.

 **NOTE:** The maximum logical drive size for all supported RAID levels (0, 1, 5, 10, and 50) is 2 TB. You can create multiple logical drives within the same physical array.

RAID Array

A RAID array is one or more logical drives controlled by the PERC.

Channel Redundant Logical Drives

When you create a logical drive, it is possible to use disks attached to different channels to implement channel redundancy, known as Channel Redundant Logical Drives. This configuration might be used for disks that reside in enclosures subject to thermal shutdown.

For more information refer to the Dell OpenManage Array Manager or Dell OpenManage Storage Management user guides located at: <http://support.dell.com>.

 **NOTE:** Channel redundancy applies only to controllers that have more than one channel and that attach to an external disk enclosure.

 **NOTE:** Make sure that the spans are in different backplanes, so that if one span fails, you do not lose the whole array.

Fault Tolerance

Fault tolerance is the capability of the subsystem to undergo a single drive failure per span without compromising data integrity, and processing capability. The RAID controller provides this support through redundant arrays in RAID levels 1, 5, 10 and 50. The system can still work properly even with a single disk failure in an array, though performance can be degraded to some extent.

 **NOTE:** RAID level 0 is not fault tolerant. If a drive in a RAID 0 array fails, the whole logical drive (all physical drives associated with the logical drive) will fail.

Fault tolerance is often associated with system availability because it allows the system to be available during the failures. However, this means it is also important for the system to be available during the repair of the problem. To make this possible, PERC 4/Di/Si and 4e/Di/Si support hot spare disks, and the auto-rebuild feature.

A hot spare is an unused physical disk that, in case of a disk failure in a redundant RAID array, can be used to rebuild the data and re-establish redundancy. After the hot spare is automatically moved into the RAID array, the data is automatically rebuilt on the hot spare drive. The RAID array continues to handle requests while the rebuild occurs.

Auto-rebuild allows a failed drive to be replaced and the data automatically rebuilt by "hot-swapping" the drive in the same drive bay. The RAID array continues to handle requests while the rebuild occurs.

Consistency Check

The Consistency Check operation verifies correctness of the data in logical drives that use RAID levels 1, 5, 10, and 50. (RAID 0 does not provide data

redundancy). For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive.

 **NOTE:** It is recommended that you perform a consistency check at least once a month.

Background Initialization

Background initialization is a consistency check that is forced when you create a logical drive. The difference between a background initialization and a consistency check is that a background initialization is forced on new logical drives. This is an automatic operation that starts 5 minutes after you create the drive.

Background initialization is a check for media errors on physical drives. It ensures that striped data segments are the same on all physical drives in an array. The background initialization rate is controlled by the rebuild rate set using the BIOS Configuration Utility. The default and recommended rate is 30%. Before you change the rebuild rate, you must stop the background initialization or the rate change will not affect the background initialization rate. After you stop background initialization and change the rebuild rate, the rate change takes effect when you restart background initialization.

Patrol Read

Patrol Read involves the review of your system for possible hard drive errors that could lead to drive failure, then action to correct errors. The goal is to protect data integrity by detecting physical drive failure before the failure can damage data. The corrective actions depend on the array configuration and type of errors.

Patrol Read starts only when the controller is idle for a defined period of time and no other background tasks are active, though it can continue to run during heavy I/O processes.

You can use the BIOS Configuration Utility to select the Patrol Read options, which you can use to set automatic or manual operation, or disable Patrol Read. Perform the following steps to select a Patrol Read option:

1. Select **Objects**—> **Adapter** from the **Management Menu**.

The **Adapter** menu displays.

2. Select **Patrol Read Options** from the **Adapter** menu.

The following options display:

- o **Patrol Read Mode**
- o **Patrol Read Status**
- o **Patrol Read Control**

3. Select **Patrol Read Mode** to display the Patrol Read mode options:
 - o **Manual** - In manual mode, you must initiate the Patrol Read.
 - o **Auto** - In auto mode, the firmware initiates the Patrol Read on a scheduled basis.
 - o **Manual Halt** - Use manual halt to stop the automatic operation, then switch to manual mode.
 - o **Disable** - Use this option to disable Patrol Read.
4. If you use **Manual** mode, perform the following steps to initiate Patrol Read:
 - a. Select **Patrol Read Control** and press <Enter>.
 - b. Select **Start** and press <Enter>.

 **NOTE:** Pause/Resume is not a valid operation when Patrol Read is set to **Manual** mode.

5. Select **Patrol Read Status** to display the number of iterations completed, the current state of the Patrol Read (active or stopped), and the schedule for the next execution of Patrol Read.

Disk Striping

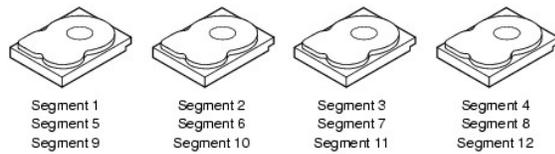
Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. PERC 4/Di/Si and 4e/Di/Si support stripe sizes of 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, and 128 KB. It is recommended that you keep stripe sizes the same across RAID arrays.

NOTE: Using a 2 KB or 4 KB stripe size is not recommended due to performance implications. Use 2 KB or 4 KB only when required by the applications used. The default stripe size is 64 KB. Do not install an operating system on a logical drive with less than a 16 KB stripe size.

For example, in a four-disk system using only disk striping (used in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple drives 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)



Stripe Width

Stripe width is the number of disks involved in an array where striping is implemented. For example, a four-disk array with disk striping has a stripe width of four.

Stripe Size

The stripe size is the length of the interleaved data segments that the RAID controller writes across multiple drives. PERC 4/Di/Si and 4e/Di/Si support stripe sizes of 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, and 128 KB.

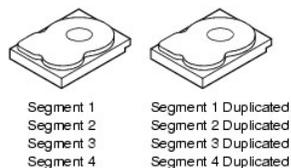
NOTE: Using a 2 KB or 4 KB stripe size is not recommended due to performance implications. Use 2 KB or 4 KB only when required by the applications used. The default stripe size is 64 KB. Do not install an operating system on a logical drive with less than a 16 KB stripe size.

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 reconstruct the failed disk. The primary advantage of disk mirroring is that it provides 100% data redundancy. Because the contents of the disk are completely written to a second disk, it does not matter if one of the disks fails. Both disks contain the same data at all times. Either drive can act as the operational drive.

Disk mirroring provides 100% redundancy, but is expensive because each drive in the system must be duplicated. [Figure 2-2](#) shows an example of disk mirroring.

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



Parity

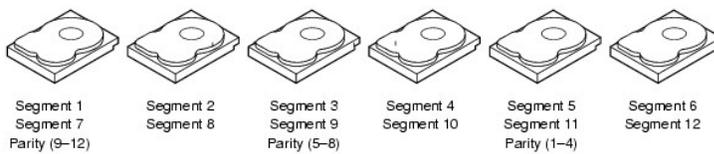
Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets. Parity data does not fully duplicate the parent data sets. In RAID, this method is applied to entire drives or stripes across all disk drives in an array. The types of parity are shown in [Table 2-1](#).

Table 2-1. Types of Parity

Parity Type	Description
Dedicated	The parity of the data on two or more disk drives is stored on an additional disk.
Distributed	The parity data is distributed across more than one drive in the system.

If a single disk drive fails, it can be rebuilt from the parity and the data on the remaining drives. RAID level 5 combines distributed parity with disk striping, as shown in [Figure 2-3](#). Parity provides redundancy for one drive failure without duplicating the contents of entire disk drives, but parity generation can slow the write process.

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



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

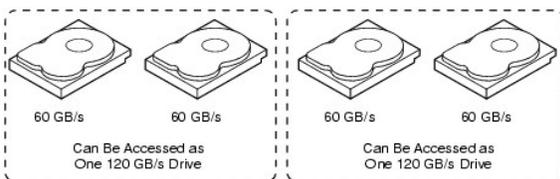
Disk Spanning

Disk spanning allows multiple physical drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 20 GB drives can be combined to appear to the operating system as a single 80 GB drive.

Spanning alone does not provide reliability or performance enhancements. Spanned logical drives must have the same stripe size and must be contiguous. In [Figure 2-4](#), RAID 1 arrays are turned into a RAID 10 array.

NOTE: Make sure that the spans in a RAID 10 array are in different backplanes, so that if one span fails, you won't lose the whole array.

Figure 2-4. Example of Disk Spanning



NOTE: Spanning two contiguous RAID 0 logical drives does not produce a new RAID level or add fault tolerance. It does increase the size of the logical volume and improves performance by doubling the number of spindles.

Spanning for RAID 10 or RAID 50

[Table 2-2](#) describes how to configure RAID 10 and RAID 50 by spanning. The PERC 4/Di/Si and 4e/Di/Si family supports spanning for RAID 1 and RAID 5 only. The logical drives must have the same stripe size and the maximum number of spans is eight. The full drive size is used when you span logical drives; you cannot specify a smaller drive size.

See [RAID Configuration and Management](#) for detailed procedures for configuring arrays and logical drives, and spanning the drives.

Table 2-2. Spanning for RAID 10 and RAID 50

Level	Description
10	Configure RAID 10 by spanning two contiguous RAID 1 logical drives. The RAID 1 logical drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 logical drives. The RAID 5 logical drives must have the same stripe size.

Hot Spares

A hot spare is an extra, unused disk drive that is part of the disk subsystem. It is usually in standby mode, ready for service if a drive fails. Hot spares permit you to replace failed drives without system shutdown or user intervention. PERC 4/Di/Si and 4e/Di/Si implement automatic and transparent rebuilds of failed drives using hot spare drives, providing a high degree of fault tolerance and zero downtime.

 **NOTE:** When running RAID 0 and RAID 5 logical drives on the same set of physical drives (a sliced configuration), a rebuild to a hot spare will not occur after a drive failure until the RAID 0 logical drive is deleted.

The PERC 4/Di/Si and 4e/Di/Si RAID management software allows you to specify physical drives as hot spares. When a hot spare is needed, the RAID controller assigns the hot spare that has a capacity closest to and at least as great as that of the failed drive to take the place of the failed drive. The failed drive is removed from the logical drive and marked ready awaiting removal once the rebuild to a hot spare begins. See [Table 4-12 in Assigning RAID Levels](#) for detailed information about the minimum and maximum number of hard drives supported by each RAID level for each RAID controller. You can make hot spares of the physical drives that are not in a RAID logical drive.

 **NOTE:** If a rebuild to a hot spare fails for any reason, the hot spare drive will be marked as "failed". If the source drive fails, both the source drive and the hot spare drive will be marked as "failed".

There are two types of hot spares:

- 1 Global Hot Spare
- 1 Dedicated Hot Spare

Global Hot Spare

A global hot spare drive can be used to replace any failed drive in a redundant array as long as its capacity is equal to or larger than the coerced capacity of the failed drive. A global hot spare defined on any channel should be available to replace a failed drive on both channels.

Dedicated Hot Spare

A dedicated hot spare can be used to replace a failed drive only in a selected array. One or more drives can be designated as member of a spare drive pool; the most suitable drive from the pool is selected for fail over. A dedicated hot spare is used before one from the global hot spare pool.

Hot spare drives can be located on any RAID channel. Standby hot spares (not being used in RAID array) are polled every 60 seconds at a minimum, and their status made available in the array management software. PERC 4/Di/Si and 4e/Di/Si offer the ability to rebuild with a disk that is in a system, but not initially set to be a hot spare.

Observe the following parameters when using hot spares:

- 1 Hot spares are used only in arrays with redundancy, for example, RAID levels 1, 5, 10, and 50.
- 1 A hot spare connected to a specific RAID controller can be used to rebuild a drive that is connected to the same controller only.
- 1 You must assign the hot spare to one or more drives through the controller's BIOS or use array management software to place it in the hot spare pool.
- 1 A hot spare must have free space equal to or greater than the drive it would replace. For example, to replace an 18 GB drive, the hot spare must be 18 GB or larger.

Disk Rebuilds

When a physical drive in a RAID array fails, you can rebuild the drive by recreating the data that was stored on the drive before it failed. The RAID controller uses hot spares to rebuild failed drives automatically and transparently, at user-defined rebuild rates. If a hot spare is available, the rebuild can start automatically when a drive fails. If a hot spare is not available, the failed drive must be replaced with a new drive so the data on the failed drive can be rebuilt. Rebuilding can be done only in arrays with data redundancy, which includes RAID 1, 5, 10, and 50.

The failed physical drive is removed from the logical drive and marked ready awaiting removal once the rebuild to a hot spare begins. If the system goes down during a rebuild, the RAID controller automatically restarts the rebuild after the system reboots.

 **NOTE:** When the rebuild to a hot spare begins, the failed drive is often removed from the logical drive before management applications, such as Dell OpenManage Storage Management, or Dell OpenManage Storage Management, detect the failed drive. When this occurs, the events logs show the drive rebuilding to the hot spare without showing the failed drive. The formerly failed drive will be marked as "ready" after a rebuild begins to a hot spare.

 **NOTE:** If a rebuild to a hot spare fails for any reason, the hot spare drive will be marked as "failed". If the source drive fails, both the source drive and the hot spare drive will be marked as "failed".

An automatic drive rebuild will not start if you replace a drive during an online capacity expansion or RAID level migration. The rebuild must be started manually after the expansion or migration procedure is complete.

Rebuild Checkpoint

The Dell PERC firmware has a feature to resume a rebuild on a physical drive in case of an abrupt power loss or if the server rebooted in the middle of a rebuild operation. In any of the following cases, however, a rebuild will not resume:

- 1 A configuration mismatch is detected on the controller.
- 1 A reconstruction is also currently in progress.
- 1 The logical drive is now owned by the peer node.

Rebuild Rate

The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives. A rebuild rate of 100 percent means the system gives priority to rebuilding the failed drives.

The rebuild rate can be configured between 0 percent and 100 percent. At 0 percent, the rebuild is done only if the system is not doing anything else. At 100 percent, the rebuild has a higher priority than any other system activity. Using 0 or 100 percent is not recommended. The default rebuild rate is 30 percent.

Hot Swap

A hot swap is the manual replacement of a defective physical disk unit while the computer is still running. When a new drive has been installed, a rebuild will occur automatically if:

- 1 The newly inserted drive is the same size as or larger than the failed drive
- 1 It is placed in the same drive bay as the failed drive it is replacing

The RAID controller can be configured to detect the new disks and rebuild the contents of the disk drive automatically.

SCSI Physical Drive States

The Physical SCSI drive states are described in [Table 2-3](#).

Table 2-3. SCSI Physical Drive States

State	Description
Online	The physical drive is working normally and is a part of a configured logical drive.
Ready	The physical drive is functioning normally but is not part of a configured logical drive and is not designated as a hot spare.

Hot Spare	The physical drive is powered up and ready for use as a spare in case an online drive fails.
Fail	A fault has occurred in the physical drive, placing it out of service.
Rebuild	The physical drive is being rebuilt with data from a failed drive.

Logical Drive States

The logical drive states are described in [Table 2-4](#).

Table 2-4. Logical Drive States

State	Description
Optimal	The logical drive operating condition is good. All configured physical drives are online.
Degraded	The logical drive operating condition is not optimal. One of the configured physical drives has failed or is offline.
Failed	The logical drive has failed.
Offline	The logical drive is not available to the RAID controller.

Enclosure Management

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

RAID Levels

The RAID controller supports RAID levels 0, 1, 5, 10, and 50. The supported RAID levels are summarized in the following section. In addition, it supports independent drives (configured as RAID 0.) The following sections describe the RAID levels in detail.

Summary of RAID Levels

RAID 0 uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.

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

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

RAID 10, a combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy, but uses a larger number of spans.

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

 **NOTE:** Running RAID 0 and RAID 5 logical arrays on the same set of physical disks (a sliced configuration) is not recommended. In the event of a disk failure, the RAID 0 logical drive will cause any rebuild attempt to fail.

Selecting a RAID Level

To ensure the best performance, you should select the optimal RAID level when you create a system drive. The optimal RAID level for your disk array depends on a number of factors:

- 1 The number of physical drives in the disk array
- 1 The capacity of the physical drives in the array
- 1 The need for data redundancy
- 1 The disk performance requirements

RAID 0

RAID 0 provides disk striping across all drives in the RAID array. RAID 0 does not provide any data redundancy, but does offer the best performance of any RAID level. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the array. The size of each block is determined by the stripe size parameter, set during the creation of the RAID set. RAID 0 offers high bandwidth.

 **NOTE:** RAID level 0 is not fault tolerant. If a drive in a RAID 0 array fails, the whole logical drive (all physical drives associated with the logical drive) will fail.

By breaking up a large file into smaller blocks, the RAID controller can use several drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. RAID 0 is also used to denote an "independent" or single drive.

[Table 2-5](#) provides an overview of RAID 0.

Table 2-5. RAID 0 Overview

Uses	Provides high data throughput, especially for large files. Any environment that does not require fault tolerance.
Strong Points	Provides increased data throughput for large files. No capacity loss penalty for parity.
Weak Points	Does not provide fault tolerance or high bandwidth. All data lost if any drive fails.
Drives	1 to 32

RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. [Table 2-6](#) provides an overview of RAID 1.

Table 2-6. RAID 1 Overview

Uses	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
Strong Points	Provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
Weak Points	Requires twice as many disk drives. Performance is impaired during drive rebuilds.
Drives	2

RAID 5

RAID 5 includes disk striping at the block level and parity. In RAID 5, the parity information is written to several drives. RAID 5 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously.

RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments.

[Table 2-7](#) provides an overview of RAID 5.

Table 2-7. RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. Use also for office automation and
-------------	--

	online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). Disk drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 28

RAID 10

RAID 10 is a combination of RAID 0 and RAID 1. RAID 10 consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks, then mirrors the blocks of data to each RAID 1 RAID set. Each RAID 1 RAID set then duplicates its data to its other drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. Up to 8 spans can be supported by RAID 10.

[Table 2-8](#) provides an overview of RAID 10.

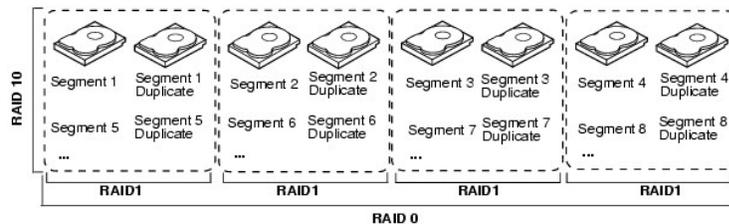
Table 2-8. RAID 10 Overview

Uses	Appropriate when used with data storage that needs 100% redundancy of mirrored arrays and that also needs the enhanced I/O performance of RAID 0 (striped arrays.) RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	2n, where n is greater than 1.

In [Figure 2-5](#), logical drive 0 is created by distributing data across four arrays (arrays 0 through 3). Spanning is used because one logical drive is defined across more than one array. Logical drives defined across multiple RAID 1 level arrays are referred to as RAID level 10, (1+0). To increase performance, by enabling access to multiple arrays simultaneously, data is striped across arrays.

Using RAID level 10, rather than a simple RAID set, up to 8 spans can be supported, and up to 8 drive failures (one failure per span) can be tolerated, though less than total disk drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 1 level array.

Figure 2-5. RAID 10 Level Logical Drive



RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 includes both parity and disk striping across multiple arrays. RAID 50 is best implemented on two RAID 5 disk arrays with data striped across both disk arrays.

RAID 50 breaks up data into smaller blocks, then stripes the blocks of data to each RAID 5 disk set. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID level 50 can support up to 8 spans and tolerate up to 8 drive failures (one failure per span), though less than total disk drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 1 level array.

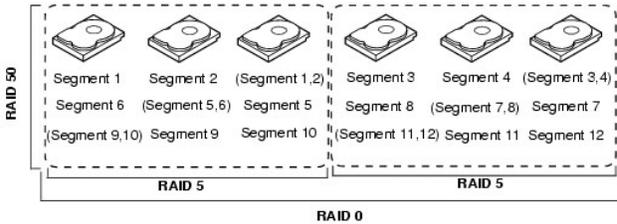
[Table 2-9](#) provides an overview of RAID 50.

Table 2-9. RAID 50 Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.
Strong Points	Provides high data throughput, data redundancy and very good performance.
Weak Points	Requires 2 to 8 times as many parity drives as RAID 5.
Drives	6 to 28 Dell supports the use of two channels with a maximum of 14 physical drives per channel.

[Figure 2-6](#) provides an example of a RAID 50 level logical drive.

Figure 2-6. RAID 50 Level Logical Drive



RAID Configuration Strategies

The most important factors in RAID array configuration are:

- 1 Logical drive availability (fault tolerance)
- 1 Logical drive performance
- 1 Logical drive capacity

You cannot configure a logical drive that optimizes all three factors, but it is easy to choose a logical drive configuration that maximizes one factor at the expense of another factor. For example, RAID 1 (mirroring) provides excellent fault tolerance, but requires a redundant drive. The following subsections describe how to use the RAID levels to maximize logical drive availability (fault tolerance), logical drive performance, and logical drive capacity.

Maximizing Fault Tolerance

Fault tolerance is achieved through the ability to perform automatic and transparent rebuilds using hot spare drives, and hot swaps. A hot spare drive is an unused online available drive that PERC 4/Di/Si and 4e/Di/Si instantly plug into the system when an active drive fails. After the hot spare is automatically moved into the RAID array, the failed drive is automatically rebuilt on the spare drive. The RAID array continues to handle requests while the rebuild occurs.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running hot swap drives. Auto-Rebuild in the BIOS Configuration Utility allows a failed drive to be replaced and automatically rebuilt by "hot-swapping" the drive in the same drive bay. The RAID array continues to handle requests while the rebuild occurs, providing a high degree of fault tolerance and zero downtime. [Table 2-10](#) describes the fault tolerance features of each RAID level.

Table 2-10. RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data lost if any drive fails. Disk striping writes data across multiple disk drives instead of just one disk drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high performance but do not require fault tolerance.
1	Provides complete data redundancy. If one disk drive fails, the contents of the other disk drive can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100% data redundancy. Since the contents of the disk drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 is ideal for any application

	that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire disk drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all disk drives in an array. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.
10	Provides complete data redundancy using striping across spanned RAID 1 arrays. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored arrays. RAID 10 can sustain a drive failure in each mirrored array and maintain drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 arrays. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information. RAID 50 can sustain one drive failure per RAID 5 array and still maintain data integrity.

Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID array appears to the host computer as a single storage unit or as multiple logical units. I/O is faster because drives can be accessed simultaneously. [Table 2-11](#) describes the performance for each RAID level.

Table 2-11. RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers the best performance of any RAID level. RAID 0 breaks up data into smaller blocks, then writes a block to each drive in the array. Disk striping writes data across multiple disk drives instead of just one disk drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.
5	RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments. Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Disk drive performance is reduced when a drive is being rebuilt. Clustering can also reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
10	RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped arrays), which provides high data transfer rates. Spanning increases the size of the logical volume and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 array.
50	RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the size of the logical volume and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 array.

Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. There are several variables to consider. Mirrored data and parity data require more storage space than striping alone (RAID 0). Parity generation uses algorithms to create redundancy and requires less space than mirroring. [Table 2-12](#) explains the effects of the RAID levels on storage capacity.

Table 2-12. RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (disk striping) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 0 provides maximum storage capacity for a given set of physical disks.
1	With RAID 1 (mirroring), data written to one disk drive is simultaneously written to another disk drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire disk drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
10	RAID 10 requires twice as many drives as all other RAID levels except RAID 1. RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity. Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	RAID 50 requires two to four times as many parity drives as RAID 5. This RAID level works best when used with data that requires medium to large capacity.

RAID Availability

RAID Availability Concept

Data availability without downtime is essential for many types of data processing and storage systems. Businesses want to avoid the financial costs and customer frustration associated with downed servers. RAID helps you maintain data availability and avoid downtime for the servers that provide that data. RAID offers several features, such as spare drives and rebuilds, that you can use to fix any hard drive problems, while keeping the server(s) running and data available. The following subsections describe these features.

Spare Drives

You can use spare drives to replace failed or defective drives in an array. A replacement drive must be at least as large as drive it replaces. Spare drives include hot swaps, hot spares, and cold swaps.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions). The backplane and enclosure must support hot swap in order for the functionality to work.

Hot spare drives are physical drives that power up along with the RAID drives and operate in a standby state. If a hard drive used in a RAID logical drive fails, a hot spare automatically takes its place and the data on the failed drive is rebuilt on the hot spare. Hot spares can be used for RAID levels 1, 5, 10, and 50.

 **NOTE:** If a rebuild to a hotspare fails for any reason, the hotspare drive will be marked as "failed". If the source drive fails, both the source drive and the hot spare drive will be marked as "failed".

A cold swap requires that you power down the system before replacing a defective hard drive in a disk subsystem.

Sector Re-assignment

Sector reassignment is done automatically by either the drive or the RAID firmware whenever a media defect is encountered.

Rebuilding

If a hard drive fails in an array that is configured as a RAID 1, 5, 10, or 50 logical drive, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed disks. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed drives are available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

Configuration Planning

Factors to consider when planning a configuration are the number of hard disk drives the RAID controller can support, the purpose of the array, and the availability of spare drives.

Each type of data stored in the disk subsystem has a different frequency of read and write activity. If you know the data access requirements, you can more successfully determine a strategy for optimizing the disk subsystem capacity, availability, and performance.

Servers that support video on demand typically read the data often, but write data infrequently. Both the read and write operations tend to be long. Data stored on a general-purpose file server involves relatively short read and write operations with relatively small files.

Number of Hard Disk Drives

Your configuration planning depends in part on the number of hard disk drives that you want to use in a RAID array. The number of drives in an array determines the RAID levels that can be supported. See [Table 4-12 in Assigning RAID Levels](#) for detailed information about the minimum and maximum number of hard drives supported by each RAID level for each RAID controller.

Array Purpose

Important factors to consider when creating RAID arrays include availability, performance, and capacity. Define the major purpose of the disk array by answering questions related to these factors, such as the following, which are followed by suggested RAID levels for each situation:

- 1 Will this disk array increase the system storage capacity for general-purpose file and print servers? Use RAID 5, 10, or 50.
- 1 Does this disk array support any software system that must be available 24 hours per day? Use RAID 1, 5, 10, or 50.
- 1 Will the information stored in this disk array contain large audio or video files that must be available on demand? Use RAID 0.
- 1 Will this disk array contain data from an imaging system? Use RAID 0 or 10.

Fill out [Table 2-13](#) to help you plan the array configuration. Rank the requirements for your array, such as storage space and data redundancy, in order of importance, then review the suggested RAID levels. Refer to [Table 4-12](#) for the minimum and maximum number of drives allowed per RAID level.

Table 2-13. Factors to Consider for Array Configuration

Requirement	Rank	Suggested RAID Level(s)
Storage space		RAID 0, RAID 5
Data redundancy		RAID 5, RAID 10, RAID 50
Hard drive performance and throughput		RAID 0, RAID 10
Hot spares (extra hard drives required)		RAID 1, RAID 5, RAID 10, RAID 50

[Back to Contents Page](#)

[Back to Contents Page](#)

Features

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [PassThru \(Legacy\) SCSI Channel](#)
- [RAID Configuration Information](#)
- [RAID Performance Features](#)
- [RAID Management Utilities](#)
- [Supported Operating Systems and Drivers](#)
- [Fault Tolerance Features](#)
- [RAID Controller Specifications](#)

This section describes the features of the RAID controller, such as the configuration features, array performance features, hardware specifications, RAID management utilities, and operating system software drivers.

Compatibility with Arrays Created on Existing RAID Controllers

The RAID controller recognizes and uses drive arrays created on existing RAID controllers without risking data loss, corruption, redundancy or configuration loss. Similarly, the arrays created on the PERC 4/Di/Si and 4e/Di/Si controllers can be transferred to other PERC 4/Di/Si and 4e/Di/Si controllers.

 **NOTE:** If you have questions about compatibility, contact your Dell Support Representative.

SMART Technology

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

Patrol Read

Patrol Read involves the review of your system for possible hard drive errors that could lead to drive failure, then action to correct the errors. The goal is to protect data integrity by detecting physical drive failure before the failure can damage data. Patrol Read adjusts the amount of RAID controller resources dedicated to Patrol Read operations based on outstanding disk I/O.

Patrol Read starts only when the controller is idle for a defined period of time and no other background tasks are active, though it can continue to run during heavy I/O processes.

You can use the BIOS Configuration Utility to select the Patrol Read options, which you can use to set automatic or manual operation, or disable Patrol Read. See [Patrol Read](#) in [RAID Configuration and Management](#) for detailed information about Patrol Read.

 **NOTE:** **Pause/Resume** is not a valid operation when Patrol Read is set to **Manual** mode.

Background Initialization

Background initialization is the automatic check for media errors on physical drives. It makes sure that striped data segments are the same on all physical drives in an array.

 **NOTE:** If you cancel background initialization, it automatically re-starts within 5 minutes. You cannot permanently cancel the background initialization.

The background initialization rate is controlled by the rebuild rate set using your array management software. The default and recommended rate is 30%. You must stop the background initialization before you change the rebuild rate or the rate change will not affect the background initialization rate. After you stop background initialization and change the rebuild rate, the rate change takes effect when you restart background initialization.

 **NOTE:** Unlike initialization of logical drives, background initialization does not clear data from the drives.

LED Operation

The LED on the drive carrier indicates the state of each drive. For internal storage enclosures, see the storage enclosure user's guide for more information about the blink patterns.

PassThru (Legacy) SCSI Channel

The RAID controller provides the ability to use a passthru (legacy) SCSI channel. This is known as "RAID/SCSI mode." Use this option for one RAID channel for hard drives and a legacy SCSI channel for removable devices or pre-existing hard drives. You can select this option in system setup interface. This is available for PERC 4/Di and 4e/Di only.

Devices attached to the SCSI channel are not controlled by the RAID firmware and function as if they are attached to a regular SCSI controller.

 **NOTE:** The passthru SCSI channel is available only on certain platforms. See your system user's guide for more information.

The devices and capabilities supported under the SCSI channel are:

- 1 Hard drives
- 1 CD drives
- 1 Tape drive units
- 1 Tape drive libraries
- 1 Support of domain validation, data CRC, double clocking, and packetization

RAID Configuration Information

[Table 3-1](#) lists the configuration features for the RAID controller.

Table 3-1. Features for RAID Configuration

Specification	PERC 4/Di/Si	PERC 4e/Di/Si
Number of logical drives and arrays supported	Up to 40 logical drives and 32 arrays per controller	Up to 40 logical drives and 32 arrays per controller
Support for hard drives with capacities of more than eight gigabytes (GB)	Yes	Yes
Online RAID level migration	Yes	Yes
Drive roaming	Yes	Yes
No reboot necessary after capacity expansion	Yes	Yes
User-specified rebuild rate	Yes	Yes

RAID Performance Features

[Table 3-2](#) displays the array performance features for the RAID controller.

Table 3-2. Array Performance Features

Specification	Description
---------------	-------------

Maximum number of scatter/gather elements	64
Drive data transfer rate	Up to 320 MB/sec
Maximum size of I/O requests	6.4 MB in 64 KB stripes
Maximum outstanding I/O requests per drive	Limited only to drive capabilities
Stripe sizes	2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB
Maximum number of concurrent commands	255 (Linux® supports only 126 concurrent commands. The 255 command limit is in the firmware and the driver limit is lower.)

RAID Management Utilities

Software utilities enable you to manage and configure the RAID system, create and manage multiple disk arrays, control and monitor multiple RAID servers, provide error statistics logging, and provide online maintenance. The utilities include:

- 1 BIOS Configuration Utility
- 1 Dell OpenManage™ Array Manager for Windows and Netware
- 1 Dell OpenManage Storage Management

BIOS Configuration Utility

The BIOS Configuration Utility configures and maintains RAID arrays, clears hard drives, and manages the RAID system. It is independent of any operating system. See [RAID Configuration and Management](#) for additional information.

Dell OpenManage Array Manager

Dell OpenManage Array Manager is used to configure and manage a storage system that is connected to a server. Array Manager runs under Novell NetWare, Windows 2000, and Windows Server 2003. Refer to the online documentation that accompanies Array Manager or the documentation section at support.dell.com for more information.

 **NOTE:** You can run the OpenManage Array Manager remotely to access NetWare, but not locally.

Dell OpenManage Storage Management

Storage Management provides enhanced features for configuring a system's locally-attached RAID and non-RAID disk storage. Storage Management enables you to perform controller and enclosure functions for all supported RAID and non-RAID controllers and enclosures from a single graphical or command-line interface without requiring use of the controller BIOS utilities. The graphical interface is wizard-driven with features for novice and advanced users and detailed online help. The command-line interface is fully-featured and scriptable.

Using Storage Management, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed drives. You can also perform data-destructive tasks. All users of Storage Management should be familiar with their storage environment and storage management.

Supported Operating Systems and Drivers

Drivers are provided to support each PERC 4e/Di/Si RAID controller for the operating systems listed in [Table 3-3](#). See [Driver Installation](#) for installation procedures for the drivers.

Table 3-3. Supported Operating Systems

Operating System	PERC 4/Di	PERC 4e/Di	PERC 4e/Si
W2K Server SP4	Y	Y	Y
W2K Advanced Server SP4	Y	Y	Y

WS 2003 Standard Server	Y	Y	Y
WS 2003 Web Server	Y	Y	Y
2003 Small Business Server (SBS)	Y	Y	Y
WS 2003 Enterprise Server	Y	Y	Y
W2K3 EM64T	N	Y	Y
RHEL v2.1 Update 3	Y	Y	Y
RHEL v3.0 Update 2 (EM64T)	N	Y	Y
RHEL v3.0 GOLD	Y	Y	Y
RHEL v3.0 Update 3 (32bit and EM64T)	N	Y	Y
RHEL 4.0 32-bit	Y	Y	Y
RHEL 4.0 EM64T	N	Y	Y
NetWare 5.1 SP8	Y	Y	Y
NetWare 6.5 SP3	Y	Y	Y

Firmware Upgrade

 **NOTICE:** Do not flash the RAID controller firmware while performing a background initialization or data consistency check as this can cause the procedure to fail.

You can download the latest firmware from the Dell website and flash it to the firmware on the board. The Dell website offers a firmware flash that can be executed from a DOS environment or one that can be launched from within the Microsoft Windows or Linux operating system. To upgrade the firmware on the RAID controller perform the following instructions:

1. Download the latest RAID controller firmware from the Dell website located at: <http://support.dell.com>.

 **NOTE:** If your system does not have a floppy disk drive, download the firmware update utility for Microsoft Windows or Linux. For systems running Novell Netware that do not have a floppy drive: create the firmware update diskette on another system, then copy the contents of the floppy to a bootable USB key or CD-ROM.

2. Use the unique instructions for each firmware update version that are posted on the Dell website to complete the firmware update process.

 **NOTE:** A reboot is required after the firmware update is complete.

Fault Tolerance Features

[Table 3-4](#) lists the features that provide fault tolerance to prevent data loss in case of a failed drive.

Table 3-4. Fault Tolerance Features

Specification	Feature
Support for SMART	Yes
Support for Patrol Read	Yes
Drive failure detection	Automatic
Drive rebuild using hot spares	Automatic
Parity generation and checking	Yes
Battery backup for NVRAM to protect configuration data	Yes
Hot-swap manual replacement of a disk unit without bringing the system down	Yes

Hot Swapping

Hot swapping is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions). The backplane and enclosure must support hot swap in order for the functionality to work.

 **NOTE:** A backplane or enclosure must support hot swapping in order for the RAID controller to support hot swapping.

Failed Drive Detection

The firmware automatically detects and rebuilds failed drives. This can be done transparently with hot spares.

RAID Controller Specifications

[Table 3-5](#) lists the RAID controller specifications.

Table 3-5. RAID Controller Specifications

Parameter	PERC 4/Di/Si	PERC 4e/Di/Si
Processor	Intel® i303 64-bit RISC processor @ 100 MHz	Intel® IOP332 I/O processor with Intel XScale Technology
Bus type	PCI Rev. 2.2	PCI Express Rev. 1.0a
PCI Express controller	Intel i303	Intel i303
Bus data transfer rate	Up to 532 MB/sec at 64/66 MHz	Up to 4 GB/sec
Cache memory size	128 MB	256 MB (DDR2)
Cache function	Write-back, Write-through, Adaptive read-ahead, Non read-ahead, Read-ahead	Write-back, Write-through, Adaptive read-ahead, Non read-ahead, Read-ahead
Flashable firmware	1 MB x 8 flash ROM	4 MB x 16 flash ROM
Nonvolatile random access memory (NVRAM)	32 KB x 8 for storing RAID configuration	32 KB x 8 for storing RAID configuration
SCSI data transfer rate	Up to 320 MB/sec per channel	Up to 320 MB/sec per channel
SCSI bus	LVD or single-ended	LVD
SCSI termination	Active	Active
Termination disable	Automatic through cable and device detection	Automatic through cable and device detection
Devices per SCSI channel	Up to 15 wide or seven narrow SCSI devices	Up to 15 wide or seven narrow SCSI devices
SCSI device types	Synchronous or asynchronous	Synchronous or asynchronous
RAID levels supported	0, 1, 5, 10, and 50	0, 1, 5, 10, and 50
Multiple logical drives/arrays per controller	Up to 40 logical drives per controller	Up to 40 logical drives per controller
Online capacity expansion	Yes	Yes
Dedicated and pool hot spare	Yes	Yes
Hot swap devices supported	Yes	Yes
Non-disk devices supported	Yes	Yes
NOTE: PERC 4/Di/Si and 4e/Di/Si do not support any non-disk devices except backplanes.		
Mixed capacity hard drives	Yes	Yes

SCSI Bus

The RAID adapter controls hard drives using Ultra320 SCSI buses (channels) over which the system transfers data in Ultra320 SCSI mode. The PERC 4/Si and 4e/Si controllers control one SCSI channel, while the PERC 4/Di and 4e/Di controls two. The SCSI channel supports up to 15 wide or seven non-wide SCSI devices at speeds up to 320 MB/sec.

SCSI Termination

The RAID controller uses active termination on the SCSI bus conforming to SCSI-3 and SCSI SPI-4 specifications. Termination enable/disable is automatic through cable detection.

SCSI Firmware

The RAID controller firmware handles all RAID and SCSI command processing and supports the features described in [Table 3-6](#).

Table 3-6. SCSI Firmware

Feature	Description
Disconnect/reconnect	Optimizes SCSI bus utilization
Tagged command queuing	Multiple tags to improve random access
Scatter/gather	Single command can transfer data to and from different memory locations
Multi-threading	Up to 189 simultaneous commands with elevator sorting and concatenation of requests per SCSI channel
Stripe size	Variable for all logical drives: 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB. The default is 64 KB. NOTE: Using a 2 KB or 4 KB stripe size is not recommended due to performance implications. Use 2 KB or 4 KB only when required by the applications used. The default stripe size is 64 KB. Do not install an operating system on a logical drive with less than a 16 KB stripe size.
Rebuild	Multiple rebuilds and consistency checks with user-definable priority.

Firmware Upgrade

You can download the latest firmware from the Dell web site and flash it to the firmware on the controller. Perform the following steps to upgrade the firmware:

1. Go to the support.dell.com web site.
2. Download the latest firmware and driver to a system that has a diskette drive.

The downloaded file is an executable that copies the firmware to a diskette.

 **NOTE:** If your system does not have a floppy disk drive, use the online flash available for Windows, NetWare, and Linux, or download the file to your hard drive and burn it to CD, then use the CD-ROM.

3. Place the diskette in the system containing the RAID controller, restart the system and boot from the diskette.
4. Run `pf1ash` to flash the firmware.

 **NOTICE:** Do not flash the firmware while performing a background initialization or data consistency check, as it can cause the procedure to fail.

Dell also provides packages for firmware update from an operating system level. Go to support.dell.com for more support.

 **NOTE:** A reboot is required after the firmware update.

RAID Management

RAID management is provided by software utilities that manage and configure the RAID system and the RAID controller, create and manage multiple disk arrays, control and monitor multiple RAID servers, provide error statistics logging, and provide online maintenance. Storage management software is included with your system. They include the following:

- 1 BIOS Configuration Utility
- 1 Dell Server Assistant®
- 1 Dell OpenManage Array Manager for Windows and Novell NetWare
- 1 Dell OpenManage Storage Management

See [RAID Configuration and Management](#) for the procedures used to manage arrays and logical drives.

[Back to Contents Page](#)

[Back to Contents Page](#)

RAID Configuration and Management

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [Entering the BIOS Configuration Utility](#)
- [Exiting the Configuration Utility](#)
- [RAID Configuration Functions](#)
- [Configuration Utility Menu](#)
- [BIOS Configuration Utility Menu Options](#)
- [Device Management](#)
- [Simple Array Setup](#)
- [Advanced Array Setup](#)
- [Managing Arrays](#)
- [Deleting Logical Drives](#)
- [Patrol Read](#)

This section describes how to configure physical drives into arrays and logical drives using the BIOS Configuration Utility. Your PERC controller can also be configured using the Dell OpenManage Array or Dell OpenManage Storage Management applications. See [RAID Management Utilities](#) in [Features](#) for information about the OpenManage applications.

Entering the BIOS Configuration Utility

The BIOS Configuration Utility configures disk arrays and logical drives. Because the utility resides in the RAID controller BIOS, its operation is independent of the operating systems on your system.

Starting the BIOS Configuration Utility

When the host computer boots, hold the <Ctrl> key and press the <M> key when a BIOS banner similar to the following displays (the text for the BIOS banner may vary slightly between controllers and BIOS versions):

```
HA -0 (Bus X Dev X) Type: PERC 4e/Di Standard FWx.xx SDRAM=xxx MB
```

```
Battery Module is Present on Adapter
```

```
1 Logical Drive found on the Host Adapter
```

```
Adapter BIOS Disabled, No Logical Drives handled by BIOS
```

```
0 Logical Drive(s) handled by BIOS
```

```
Press <Ctrl><M> to Enable BIOS
```

For each controller in the host system, the firmware version, dynamic random access memory (DRAM) size, and the status of logical drives on that controller display. After you press a key to continue, the **Management Menu** screen displays.

 **NOTE:** In the BIOS Configuration Utility, pressing <Ctrl><M> has the same effect as pressing <Enter>.

 **NOTE:** You can access multiple controllers through the BIOS Configuration Utility. Be sure to verify which controller you are currently set to edit.

Exiting the Configuration Utility

1. Press <Esc> when the **Management Menu** displays.
 2. Select **Yes** at the prompt.
 3. Reboot the system.
-

RAID Configuration Functions

 **NOTE:** The OpenManage™ Array Manager and Dell OpenManage Storage Management can perform many of the same tasks as the BIOS Configuration Utility.

After you have attached all physical drives, use a configuration utility to prepare a logical drive. Your SCSI hard drives must be organized into logical drives in an array and be able to support the RAID level that you select. If the operating system is not yet installed, use the BIOS Configuration Utility to perform this procedure. If the operating system is installed, you can use OpenManage Array Manager (for Windows and NetWare) or Dell OpenManage Storage Management.

Use the configuration utilities to perform the following tasks:

1. Configure physical arrays and logical drives.
1. Create hot spare drives.
1. Initialize one or more logical drives.
1. Access controllers, logical drives, and, physical drives individually.
1. Rebuild failed hard drives.
1. Verify that the redundancy data in logical drives using RAID level 1, 5, 10, or 50 is correct.
1. Reconstruct logical drives after changing RAID levels or adding a hard drive to an array.
1. Select a host controller to work on.

The following sections describe the menu options and provide detailed instructions used to perform the configuration tasks. The following is a list of the procedures used to configure hard disk drives into arrays and logical drives. They apply to the BIOS Configuration Utility, OpenManage Array Manager, and Dell OpenManage Storage Management. The following is a list of the configuration steps:

1. Designate hot spares (optional).

See [Designating Drives as Hot Spares](#) in this section for more information.

2. Select a configuration method.

See [Configure Menu](#) for more information.

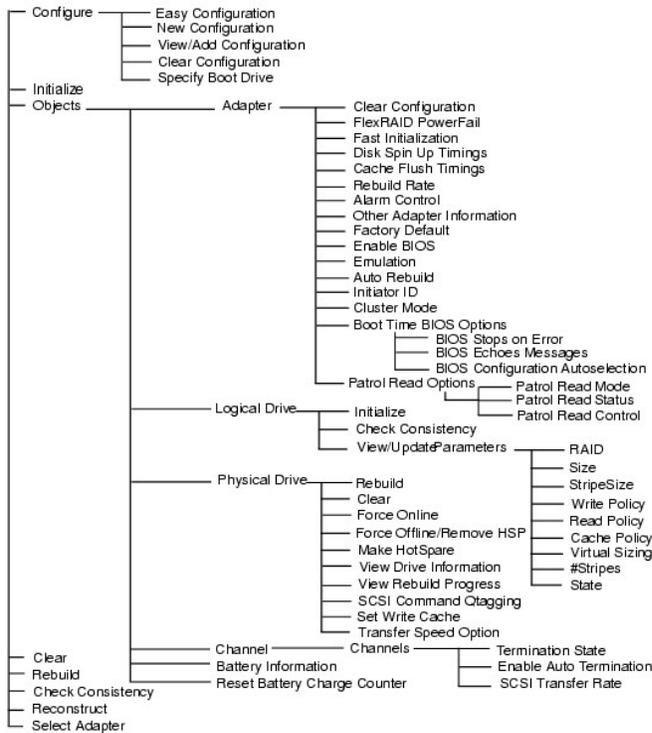
3. Create arrays using the available physical drives.
4. Define logical drives using the arrays.
5. Save the configuration information.
6. Initialize the logical drives.

See [Simple Array Setup](#) and [Advanced Array Setup](#) for the detailed configuration procedures.

Configuration Utility Menu

[Figure 4-1](#) shows the menu tree for the BIOS Configuration Utility. The following sections describe each menu item.

Figure 4-1. BIOS Configuration Utility Menu Tree



BIOS Configuration Utility Menu Options

[Table 4-1](#) describes the options for the BIOS Configuration Utility **Management Menu**. The menu and sub-menu options are explained in the following sections.

Table 4-1. BIOS Configuration Utility Menu Options

Option	Description
Configure	Select this option to configure hard disk drives into arrays and logical drives.
Initialize	Select this option to initialize one or more logical drives.
Objects	Select this option to individually access controllers, logical drives, and physical drives.
Clear	Select this option to clear the data from SCSI drives.
Rebuild	Select this option to rebuild failed hard disk drives.
Check Consistency	Select this option to verify that the redundancy data in logical drives using RAID level 1, 5, 10, or 50 is correct.
Reconstruct	Select this option to perform RAID level migration or online capacity expansion
Select Adapter	Select this option to list the adapters and select the adapter that you want to configure. The number of the selected adapter and model information display.

Configure Menu

Select **Configure** to select a method for configuring arrays and logical drives. [Table 4-2](#) displays the configuration methods, clear configuration option, and boot drive option.

Table 4-2. Configuration Menu Options

Option	Description
Easy	Select this method to perform a logical drive configuration where every physical array you define is automatically associated with exactly

Configuration	one logical drive.
New Configuration	Select this method to discard the existing configuration information and to configure new arrays and logical drives. In addition to providing the basic logical drive configuration functions, New Configuration allows you to associate logical drives with multiple arrays (spanning.)
View/Add Configuration	Select this method to examine the existing configuration and/or to specify additional arrays and logical drives. View/Add Configuration provides the same functions available in New Configuration .
Clear Configuration	Select this option to erase the current configuration information from the non-volatile memory on the RAID controller.
Specify Boot Drive	Select this option to specify a logical drive as the boot drive on this adapter.

Initialize Menu

 **NOTE:** See [Simple Array Setup](#) or [Advanced Array Setup](#) for steps for initializing logical drives.

Select **Initialize** from the **Management Menu** to initialize one or more logical drives. Press the space bar to select a single drive or the <F2> key to select all drives for initialization. This action typically follows the configuration of a new logical drive.

 **NOTICE:** Initializing a logical drive destroys all data on the logical drive.

Objects Menu

Select **Objects** from the **Management Menu** to access the adapters, logical drives, physical drives, and SCSI channels individually. You can also change settings for each object. The **Objects** menu options are described in the following sections.

Adapter

Select **Objects**→ **Adapter** to select a controller (if the computer has more than one) and to modify parameters. [Table 4-3](#) describes the **Adapter** menu options.

Table 4-3. Adapter Menu Options

Option	Description
Clear Configuration	Select this option to erase the current configuration from the controller non-volatile memory.
FlexRAID PowerFail	Select this option to enable or disable the FlexRAID PowerFail feature. This option allows drive reconstruction, rebuild, and check consistency to continue when the system restarts because of a power failure, reset, or hard boot.
Fast Initialization	Select this option to write zeros to the first sector of the logical drive so that initialization occurs in 2 – 3 seconds. When this option is set to Disabled , a full initialization takes place on the entire logical drive. On a larger array (over 5 arrays), it is best to set fast initialization to Disabled , then initialize. Otherwise, the controller will run a background consistency check within five minutes of reboot or RAID 5 creation.
Disk Spin up Timings	Select this option to set the method and timing for spinning up the hard drives.
Cache Flush Timings	Select this option to set the cache flush interval to once every 2, 4, 6, 8, or 10 seconds. The default is 4.
Rebuild Rate	Use this option to select the rebuild rate for drives attached to the selected adapter. The rebuild rate is the percentage of the system resources dedicated to rebuilding a failed drive. A rebuild rate of 100 percent means the system is totally dedicated to rebuilding the failed drive. The default is 30 percent.
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator. The alarm sounds when there is a change in a drive state, such as when a drive fails or when a rebuild is complete.
Other Adapter Information	Provides general information about the adapter, such as the firmware version, and BIOS version.
Factory Default	Select this option to load the default BIOS Configuration Utility settings.
Enable BIOS	Select this option to enable or disable the BIOS on the adapter. If the boot device is on the RAID controller, the BIOS must be enabled; otherwise, the BIOS should be disabled or it might not be possible to use a boot device elsewhere.
Emulation	You can operate in the I2O mode or mass storage mode Dell recommends that you use only mass storage mode, and Dell drivers only.
Auto Rebuild	Set to Enabled to automatically rebuild drives when they fail.
Initiator ID	Displays the Initiator ID for the cluster card. It cannot have the same ID as the other node. The default is 7.
Boot Time BIOS Options	Use this to select the following options for BIOS actions during bootup: BIOS Stops on Error: When set to On, the BIOS stops in case of a problem with the configuration. This gives you the option to enter the configuration utility to resolve the problem. The default is on.

	<p>BIOS Echoes Messages: When set to On (the default), all controller BIOS messages display during bootup.</p> <p>BIOS Configuration Autoselection: Use this option if there is a mismatch between configuration data in the drives and NVRAM, so you can select a method to resolve it. The options are <code>NVRAM</code>, <code>Disk</code>, and <code>User</code>. The default is <code>User</code>.</p>
Patrol Read Options	<p>Patrol Read involves the review of your system for possible hard drive errors that could lead to drive failure, then action to correct errors. The goal is to protect data integrity by detecting physical drive failure before the failure can damage data. Patrol Read occurs only when the controller is idle for a defined period of time and no other background tasks are active.</p> <p>Patrol Read Options gives you the ability to start and stop Patrol Reads, display Patrol Read status, and set the Patrol Read mode. See Patrol Read for detailed information about the Patrol Read.</p>

Patrol Read Options

[Table 4-4](#) describes the **Patrol Read Options** submenu. See [Patrol Read](#) for detailed information about the Patrol Read.

Table 4-4. Patrol Read Options Menu

Option	Description
Patrol Read Mode	Use this option to set Patrol Read for manual operation (user-initiated) or automatic operation, or to disable Patrol Read.
Patrol Read Status	Displays the number of iterations completed, the current state of the Patrol Read (active or stopped), and the schedule for the next execution of Patrol Read.
Patrol Read Control	Use this option to start or stop Patrol Read.

Logical Drive

Select **Objects**→ **Logical Drive** to select a logical drive and to perform the actions listed in [Table 4-5](#).

Table 4-5. Logical Drive Menu Options

Option	Description
Initialize	Initializes the selected logical drive. Do this for every logical drive that is configured.
Check Consistency	Verifies the correctness of the redundancy data in the selected logical drive. This option is available only if RAID level 1, 5, 10, or 50 is used. The RAID controller automatically corrects any differences found in the data.
View/Update Parameters	Displays the properties of the selected logical drive. You can modify the cache write policy, read policy, and the input/output (I/O) policy from this menu.

Physical Drive

Select **Objects**→ **Physical Drive** to select a physical device and to perform the operations listed in the table below. The physical drives in the computer are listed. Move the cursor to the desired device and press <Enter> to display the screen.

[Table 4-6](#) displays the operations you can perform on the physical drives.

Table 4-6. Physical Drive Menu Options

Option	Description
Rebuild	Rebuilds the selected physical drive.
Rebuild	Select this option to rebuild a failed hard disk drive.
Clear	Select this option to clear the data from SCSI drives.
Force Online	Changes the state of the selected hard drive to online.
Force Offline/Remove HSP	Changes the state of the selected hard drive to offline.
Make HotSpare	Designates the selected hard drive as a hot spare.
View Drive Information	Displays the drive properties for the selected physical device.

View Rebuild Progress	Indicates how much of the rebuild has been completed.
Set Write Cache	Select this option to enable or disable write cache on this device. See Logical Drive Parameters and Descriptions in this section for more information about write cache policy.
Transfer Speed Option	Selects the speed at which data is transferred. Displays a menu that contains the options Negotiation=Wide , and Set Transfer Speed . The maximum transfer speed is 320 M .

Channel

Select **Objects**→**Channel** to select a SCSI channel on the currently selected controller. After you select a channel, press <Enter> to display the options for that channel. [Table 4-7](#) describes the SCSI channel menu options.

Table 4-7. SCSI Channel Menu Options

Option	Description
Termination State	When set to enabled, the RAID controller is terminated. When set to disabled, it is not terminated. Normally, you do not need to change this setting; the RAID controller automatically sets this option.
Enable Auto Termination	Select this option to enable or disable auto termination of the SCSI bus.
SCSI Transfer Rate	Used to select the SCSI transfer rate. The options are Fast , Ultra , Ultra-2 , and 160M . NOTE: The disk transfer rate is set for each disk, while the SCSI channel transfer rate controls the speed of the bus. No matter how fast you set the disk transfer rate, the speed depends on the SCSI channel transfer rate.

Clear Menu

You can clear the data from SCSI drives using the configuration utilities. See [Clearing Physical Drives](#) for more information and the procedure for clearing the data.

Rebuild Menu

Select **Rebuild** from the **Management Menu** to rebuild one or more failed physical drives. See [Rebuilding Failed Hard Drives](#) for more information and the procedure to perform a drive rebuild.

Check Consistency Menu

Select **Check Consistency** to verify the redundancy data in logical drives that use RAID levels 1, 5, 10, and 50. See [Checking Data Consistency](#) for more information and the procedure to perform a check consistency.

Reconstruct Menu

Select **Reconstruct** to change the RAID level of an array or add a physical drive to an existing array. RAID level migration changes the array from one RAID level to another. Online capacity expansion is the addition of hard disk drives to increase storage capacity.

Device Management

Device Management Functions

This section deals with device management, which means management of the physical devices. This includes the physical drives, hot spares, drive migration, and drive roaming. See [Drive Roaming](#) and [Drive Migration](#) for details about these procedures.

Physical Drive Selection Menu

The configuration utility offers the **Physical Drive Selection Menu** which you can use to perform actions on the physical drives in an array, such as rebuilding a drive or making a hot spare online or offline. Some of these actions are described in detail in other sections of this chapter. Perform the following steps to view the actions you can select.

1. On the **Management Menu** select **Objects--> Physical Drive**.

A physical drive selection screen appears.

2. Select a hard drive in the **READY** state and press <Enter> to display the action menu for the physical drives.

The menu items are:

- 1 Rebuild
- 1 Clear
- 1 Force Online
- 1 Force Offline/Remove HSP
- 1 Make HotSpare
- 1 View Drive Information
- 1 View Rebuild Progress
- 1 SCSI Command Qtagging
- 1 Set Write Cache
- 1 Transfer Speed Option

Device Configuration

You can fill in the following table to list the devices assigned to Channel 1. The PERC 4/Si and 4e/Si controllers have one channel, and the PERC 4/Di and 4e/Di have two.

Use [Table 4-8](#) to list the devices that you assign to each SCSI ID for SCSI Channel 1.

Table 4-8. Configuration for SCSI Channel 1

SCSI Channel 1	
SCSI ID	Device Description
0	
1	
2	
3	
4	
5	
6	
7	Reserved for host controller.
8	
9	
10	
11	
12	
13	
14	
15	

Simple Array Setup

This section describes the steps used in **Easy Configuration** to set up a simple array and create logical drives. In **Easy Configuration**, each physical array you create is associated with exactly one logical drive, so you cannot span arrays. In addition, in **Easy Configuration**, you cannot change the logical drive size.

You can modify the following logical drive parameters, which are described in [Table 4-9](#). The spanning option is also described in [Table 4-9](#), though you cannot span arrays using **Easy Configuration**.

- 1 RAID level
- 1 Stripe size
- 1 Write policy
- 1 Read policy
- 1 Cache policy

Table 4-9. Logical Drive Parameters and Descriptions

Parameter	Description
RAID Level	The number of physical drives in a specific array determines the RAID levels that can be implemented with the array.
Stripe size	<p>Stripe Size specifies the size of the segments written to each drive in a RAID 1, 5, or 10 logical drive. You can set the stripe size to 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, or 128 KB. The default and recommended rate is 64 KB.</p> <p>NOTE: Using a 2 KB or 4 KB stripe size is not recommended due to performance implications. Use 2 KB or 4 KB only when required by the applications used. The default stripe size is 64 KB. Do not install an operating system on a logical drive with less than a 16 KB stripe size.</p> <p>A larger stripe size provides better read performance, especially if your computer does mostly sequential reads. However, if you are sure that your computer does random read requests more often, select a small stripe size.</p>
Write Policy	<p>Write Policy specifies the cache 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. This setting is recommended in standard mode.</p> <p>NOTICE: If WriteBack is enabled and the system is quickly turned off and on, the RAID controller may pause as the system flushes cache memory. Controllers that contain a battery backup will default to WriteBack caching.</p> <p>In Write-through caching, the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction.</p> <p>Write-through caching has a data security advantage over write-back caching. Write-back caching has a performance advantage over write-through caching.</p> <p>NOTE: Enabling clustering turns off write cache.</p>
Read Policy	<p>Read-ahead enables the read-ahead feature for the logical drive. You can set this parameter to Read-Ahead, No-Read-ahead, or Adaptive. The default is Adaptive.</p> <p>Read-ahead specifies that the controller uses read-ahead for the current logical drive. Read-ahead capability allows the adapter to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the data will be needed soon. Read-ahead supplies sequential data faster, but is not as effective when accessing random data.</p> <p>No-Read-Ahead specifies that the controller does not use read-ahead for the current logical drive.</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>
Cache Policy	<p>Cache Policy applies to reads on a specific logical drive. It does not affect the Read-ahead cache. The default is Direct I/O.</p> <p>Cached I/O specifies that all reads are buffered in cache memory.</p> <p>Direct I/O specifies that reads are not buffered in cache memory. Direct I/O does not override the cache policy settings. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory.</p>
Span	<p>The choices are:</p> <p>Yes—Array spanning is enabled for the current logical drive. The logical drive can occupy space in more than one array.</p> <p>No—Array spanning is disabled for the current logical drive. The logical drive can occupy space in only one array.</p> <p>The RAID controller supports spanning of RAID 1 and 5 arrays. You can span two or more RAID 1 arrays into a RAID 10 array and two or more RAID 5 arrays into a RAID 50 array. The maximum number of spans is eight.</p> <p>For two arrays to be spanned, they must have the same stripe width (they must contain the same number of physical drives).</p>

If logical drives have already been configured when you select **Easy Configuration**, the configuration information is not disturbed. Perform the following steps to create arrays and logical drives using **Easy Configuration**.

1. Select **Configure**—> **Easy Configuration** from the **Management Menu**.

Hot key information displays at the bottom of the screen.

2. Press the arrow keys to highlight specific physical drives.
3. Press the spacebar to associate the selected physical drive with the current array.

The selected drive changes from **READY** to **ONLIN A**[array number]-[drive number]. For example, **ONLIN A02-03** means array 2 with hard drive 3.

4. Add physical drives to the current array as desired.

Try to use drives of the same capacity in a specific array. If you use drives with different capacities in an array, all drives in the array are treated as if they have the capacity of the smallest drive in the array.

 **NOTE:** When you create a logical drive, you can select more than 2 TB of physical hard drive space, but 2 TB is the largest logical drive you can create. After you select the physical drives, you are prompted to press <Enter> to accept the 2TB logical drive size. You are then prompted to accept the next logical drive, which will be the remaining amount of physical hard drive space.

5. Press <Enter> after you finish creating the current array.

The **Select Configurable Array(s)** window appears. It displays the array and array number, such as **A-00**.

6. Press the spacebar to select the array.

 **NOTE:** You can press <F2> to display the number of drives in the array, their channel and ID, and press <F3> to display array information, such as the stripes, slots, and free space.

7. Press <F4> to add a hotspare, if desired, and select **Yes** at the prompt.

See [Designating Drives as Hot Spares](#) for more information.

8. Press <F10> to configure a logical drive.

The window at the top of the screen shows the logical drive that is currently being configured.

9. Highlight **RAID** and press <Enter> to set the RAID level for the logical drive.

The available RAID levels for the current logical drive display.

10. Select a RAID level and press <Enter> to confirm.
11. Click **Advanced Menu** to open the menu for logical drive settings.
12. Set the **Stripe Size**.
13. Set the **Write Policy**.
14. Set the **Read Policy**.
15. Set the **Cache Policy**.
16. Press <Esc> to exit the **Advanced Menu**.
17. After you define the current logical drive, select **Accept** and press <Enter>.

The array selection screen appears if any unconfigured hard drives remain.

18. Repeat [step 2](#) through [step 17](#) to configure another array and logical drive.

The RAID controller supports up to 40 logical drives per controller.

19. When finished configuring logical drives, press <Esc> to exit **Easy Configuration**.

A list of the currently configured logical drives appears.

20. Respond to the **Save** prompt.

After you respond to the prompt, the **Configure** menu appears.

21. Press <Esc> to return to the **Management Menu**.

The logical drives you configured need to be initialized to prepare them for use.

22. Select **Initialize** on the **Management Menu**.

The configured logical drives display.

 **NOTE:** When the **Fast Initialization** option in the **Objects→ Adapter** menu is set to **Disabled**, a full initialization takes place on the entire logical drive. On a larger array (over 5 arrays), it is best to set fast initialization to **Disabled**, then initialize. Otherwise, the controller will run a background consistency check within five minutes of reboot or RAID 5 creation.

 **NOTE:** A full initialization will not resume after a power loss; it will start completely over.

23. Use the arrow key to highlight a logical drive, then press the spacebar to select a logical drive or press<F2> to select all the logical drives.
24. Press <F10> to initialize the selected logical drive(s) and select **Yes** at the prompt.

A progress bar displays.

25. When the initialization is complete, press <Esc> to return to the **Management Menu**.
-

Advanced Array Setup

The following procedures describe more advanced array and logical drive setups. The difference between simple setup and advanced setup is that you can select drive size and span arrays in advanced setup. The configuration utilities offer **New Configuration** and **View/Add Configuration** options, which are described in the following procedures.

Using New Configuration

If you select **New Configuration**, the existing configuration information on the selected controller is **destroyed when the new configuration is saved**. In **New Configuration**, you can modify the following logical drive parameters:

- 1 RAID level
- 1 Logical drive size
- 1 Stripe size
- 1 Write policy
- 1 Read policy
- 1 Cache policy
- 1 Spanning of arrays

 **NOTICE:** Selecting **New Configuration** erases the existing configuration information on the selected controller. To use the existing configuration, use **View/Add Configuration**.

1. Select **Configure→ New Configuration** from the **Management Menu**.

Hot key information appears at the bottom of the screen.

2. Press the arrow keys to highlight specific physical drives.
3. Press the spacebar to associate the selected physical drive with the current array.

The selected drive changes from **READY** to **ONLINE A[*array number*]-[*drive number*]**. For example, **ONLINE A02-03** means array 2 with hard drive 3.

4. Add physical drives to the current array as desired.

Try to use drives of the same capacity in a specific array. If you use drives with different capacities in an array, all drives in the array are treated as if they have the capacity of the smallest drive in the array.

 **NOTE:** When you create a logical drive, you can select more than 2 TB of physical hard drive space, but 2 TB is the largest logical drive you can create. After you select the physical drives, you are prompted to press <Enter> to accept the 2TB logical drive size. You are then prompted to accept the next logical drive, which will be the remaining amount of physical hard drive space.

5. Press <Enter> twice after you finish creating the current array.

The **Select Configurable Array(s)** window appears. It displays the array and array number, such as **A-00**.

6. Press the spacebar to select the array.

Span information displays in the array box. You can create multiple arrays, then select them to span them.

 **NOTE:** You can press <F2> to display the number of drives in the array, their channel and ID, and <F3> to display array information, such as the stripes, slots, and free space.

7. Repeat [step 2](#) through [step 6](#) to create another array or go to [step 8](#) to configure a logical drive.
8. Press <F10> to configure a logical drive.

The logical drive configuration screen appears. **Span=Yes** displays on this screen if you select two or more arrays to span.

The window at the top of the screen shows the logical drive that is currently being configured as well as any existing logical drives.

9. Highlight **RAID** and press <Enter> to set the RAID level for the logical drive.

A list of the available RAID levels for the current logical drive appears.

10. Select a RAID level and press <Enter> to confirm.
11. Highlight **Span** and press <Enter>.

 **NOTE:** Make sure that the spans are in different backplanes, so that if one span fails, you won't lose the whole array.

12. Highlight a spanning option and press <Enter>.
13. Move the cursor to **Size** and press <Enter> to set the logical drive size.

By default, the logical drive size is set to all available space in the array(s) being associated with the current logical drive, accounting for the **Span** setting.

14. Click **Advanced Menu** to open the menu for logical drive settings.
15. Set the **Stripe Size**.
16. Set the **Write Policy**.
17. Set the **Read Policy**.
18. Set the **Cache Policy**.
19. Press <Esc> to exit the **Advanced Menu**.
20. After you define the current logical drive, select **Accept** and press <Enter>.

If space remains in the arrays, the next logical drive to be configured appears. If the array space has been used, a list of the existing logical drives appears.

21. Press any key to continue, then respond to the **Save** prompt.
22. Press <Esc> to return to the **Management Menu**.

The logical drives you configured need to be initialized to prepare them for use.

23. Select **Initialize** on the **Management Menu**.

The configured logical drives display.

 **NOTE:** A full initialization will not resume after a power loss; it will start completely over.

24. Use the arrow key to highlight a logical drive, then press the spacebar to select a logical drive or press <F2> to select all the logical drives.
25. Press <F10> to initialize the selected logical drive(s) and select **Yes** at the prompt.

A progress bar displays.

26. When the initialization is complete, press <Esc> to return to the **Management Menu**.

Using View/Add Configuration

View/Add Configuration allows you to control the same logical drive parameters as **New Configuration** *without disturbing* the existing configuration information. In addition, you can enable the **Configuration on Disk** feature.

1. Select **Configure**—> **View/Add Configuration** from the **Management Menu**.

Hot key information appears at the bottom of the screen.

2. Press the arrow keys to highlight specific physical drives.
3. Press the spacebar to associate the selected physical drive with the current array.

The selected drive changes from **READY** to **ONLIN A[*array number*]-[*drive number*]**. For example, **ONLIN A02-03** means array 2 with hard drive 3.

4. Add physical drives to the current array as desired.

Try to use drives of the same capacity in a specific array. If you use drives with different capacities in an array, all drives in the array are treated as if they have the capacity of the smallest drive in the array.

 **NOTE:** When you create a logical drive, you can select more than 2 TB of physical hard drive space, but 2 TB is the largest logical drive you can create. After you select the physical drives, you are prompted to press <Enter> to accept the 2TB logical drive size. You are then prompted to accept the next logical drive, which will be the remaining amount of physical hard drive space.

5. Press <Enter> twice after you finish creating the current array.

The **Select Configurable Array(s)** window appears. It displays the array and array number, such as **A-00**.

6. Press the spacebar to select the array.

Span information, such as **Span-1**, displays in the array box. You can create multiple arrays, then select them to span them.

 **NOTE:** You can press <F2> to display the number of drives in the array, their channel and ID, and <F3> to display array information, such as the stripes, slots, and free space.

7. Press <F10> to configure a logical drive.

The logical drive configuration screen appears. **Span=Yes** displays on this screen if you select two or more arrays to span.

8. Highlight **RAID** and press <Enter> to set the RAID level for the logical drive.

The available RAID levels for the current logical drive appear.

9. Select a RAID level and press <Enter> to confirm.
10. Highlight **Span** and press <Enter>.
11. Highlight a spanning option and press <Enter>.

The maximum number of spans is eight.

12. Move the cursor to **Size** and press <Enter> to set the logical drive size.

By default, the logical drive size is set to all available space in the array(s) associated with the current logical drive, accounting for the **Span** setting.

13. Highlight **Span** and press <Enter>.
14. Highlight a spanning option and press <Enter>.

 **NOTE:** The full drive size is used when you span logical drives; you cannot specify a smaller drive size.

15. Click **Advanced Menu** to open the menu for logical drive settings.
16. Set the **Stripe Size**.
17. Set the **Write Policy**.
18. Set the **Read Policy**.
19. Set the **Cache Policy**.
20. Press <Esc> to exit the **Advanced Menu**.
21. After you define the current logical drive, select **Accept** and press <Enter>.

If space remains in the arrays, the next logical drive to be configured appears.

22. Repeat [step 2](#) to [step 21](#) to create an array and configure another logical drive.

If all array space is used, a list of the existing logical drives appears.

23. Press any key to continue, then respond to the **Save** prompt.
24. Press <Esc> to return to the **Management Menu**.

The logical drives you configured need to be initialized to prepare them for use.

25. Select **Initialize** on the **Management Menu**.

The configured logical drives display.

 **NOTE:** A full initialization will not resume after a power loss; it will start completely over.

26. Use the arrow key to highlight a logical drive, then press the spacebar to select a logical drive or press <F2> to select all the logical drives.
27. Press <F10> to initialize the selected logical drive(s) and select **Yes** at the prompt.

A progress bar displays.

28. When the initialization is complete, press <Esc> to return to the **Management Menu**.

Managing Arrays

Your SCSI hard drives must be organized into logical drives in an array and must be able to support the RAID level that you select. This section describes:

- 1 Guidelines for connecting and configuring SCSI devices in a RAID array
- 1 Storage space in RAID 1 and RAID 5 arrays with hard disk drives of different sizes

- 1 Maximum number of hard disk drives that you can use in each RAID level
- 1 Array configuration
- 1 Logical drive properties
- 1 Clearing physical drives
- 1 Designating physical drives as hot spares
- 1 Rebuilding failed physical drives
- 1 Checking data consistency
- 1 Reconstructing logical drives
- 1 Performing an online capacity expansion
- 1 Performing drive roaming or drive migration

Guidelines for SCSI Devices in a RAID Array

Observe the following guidelines when connecting and configuring SCSI devices in a RAID array:

- 1 Consider the number of hard disk drives in the array when deciding on the RAID level to use. See [RAID Levels](#) for the number of drives supported for each array level.
- 1 Use drives of the same size and speed to maximize the effectiveness of the controller.
- 1 When replacing a failed drive in a redundant array, make sure that the replacement drive has the same or larger capacity than the smallest drive in the array (RAID 1, 5, 10, and 50).

When implementing RAID 1 or RAID 5, disk space is spanned to create the stripes and mirrors. The span size can vary to accommodate the different disk sizes. There is, however, the possibility that a portion of the largest disk in the array will be unusable, resulting in wasted disk space. For example, consider a RAID 1 array that has the following disks, as shown in [Table 4-10](#).

Table 4-10. Storage Space in a RAID 1 Array

Disk	Disk Size	Storage Space Used in Logical Drive for RAID 1 Array	Storage Space Left Unused
A	20 GB	20 GB	0 GB
B	30 GB	20 GB	10 GB

In the RAID 1 example, data is mirrored across the two disks until 20 GB on Disk A and B are completely full. This leaves 10 GB of disk space on Disk B. Data cannot be written to this remaining disk space, as there is no corresponding disk space available in the array to create redundant data.

[Table 4-11](#) provides an example of a RAID 5 array.

Table 4-11. Storage Space in a RAID 5 Array

Disk	Disk Size	Storage Space Used in Logical Drive for RAID 5 Array	Storage Space Left Unused
A	40 GB	40 GB	0 GB
B	40 GB	40 GB	0 GB
C	60 GB	40 GB	20 GB

In the RAID 5 example, data is striped across the disks until 40 GB on Disks A, B, and C are completely full. This leaves 20 GB of disk space on Disk C. Data cannot be written to this remaining disk space, as there is no corresponding disk space available in the array to create redundant data.

RAID levels 10 and 50 span RAID 1 and RAID 5 arrays, respectively. When one array fills its available storage space, the other array(s) may have additional storage space available. You can still use fill the additional available space in the larger array(s). Because there is additional storage space in the larger array(s), you can use arrays of different sizes without having to leave storage space unused. See [Storage in RAID 10 and RAID 50 Arrays](#) for more information about storage space in RAID 10 and 50 arrays.

Assigning RAID Levels

Only one RAID level can be assigned to each logical drive. [Table 4-12](#) shows the minimum and maximum number of drives required for each RAID level.

Table 4-12. Physical Drives Required for Each RAID Level

RAID Level	Minimum # of Physical Drives	Maximum # of Physical Drives for PERC 4//si and 4e/Si	Maximum # of Physical Drives for PERC 4/Di and 4e/Di
0	1	14	28
1	2	2	2
5	3	14	28
10	4	14	28
50	6	14	28

Array Configuration

After you configure and initialize the hard drives, you are ready to configure arrays. The number of drives in an array determines the RAID levels that can be supported. For information about the number of drives required for different RAID levels, see [Table 4-12](#) in [Assigning RAID Levels](#).

Logical Drives

Logical drives, also known as virtual disks, are arrays or spanned arrays that are available to the operating system. The storage space in a logical drive is spread across all the physical drives in the array or spanned arrays.

You must create one or more logical drives for each array, and the logical drive capacity must include all of the drive space in an array. You can make the logical drive capacity larger by spanning arrays. In an array of drives with mixed sizes, the smallest common drive size is used and the space in larger drives is not used. The RAID controller supports up to 40 logical drives.

Configuring Logical Drives

After you have attached all physical drives, perform the following steps to prepare a logical drive. If the operating system is not yet installed, use the BIOS Configuration Utility to perform this procedure.

1. Start the system.
2. Run your array management software.
3. Select the option to customize the RAID array.

In the BIOS Configuration Utility, use either **Easy Configuration** or **New Configuration** to customize the RAID array.

 **CAUTION:** If you select **New Configuration**, all previous configuration information will be deleted.

4. Create and configure one or more system drives (logical drives).
5. Select the RAID level, cache policy, read policy, and write policy.

 **NOTE:** Refer to the section [Summary of RAID Levels](#) for RAID level explanations and [Table 4-9](#) for information about the policy settings.

6. Save the configuration.
7. Initialize the system drives.

After initialization, you can install the operating system.

 **NOTE:** A full initialization will not resume after a power loss; it will start completely over.

See [Simple Array Setup](#) and [Advanced Array Setup](#) for detailed configuration instructions.

Spanned Drives

You can arrange arrays sequentially with an identical number of drives so that the drives in the different arrays are spanned. Spanned drives can be treated as one large drive. Data can be striped across multiple arrays as one logical drive. The maximum number of spans is eight.

You can create spanned drives using your array management software, which is the BIOS Configuration Utility.

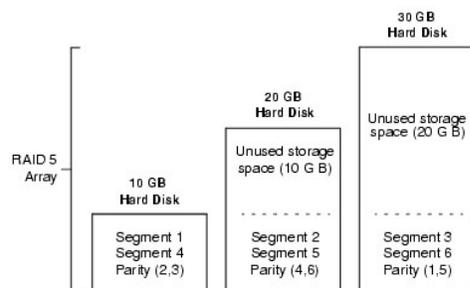
Storage in an Array with Drives of Different Sizes

For RAID levels 0 and 5, data is striped across the disks. If the hard drives in an array are not the same size, data is striped across all the drives until one or more of the drives is full. After one or more drives are full, disk space left on the other disks cannot be used. Data cannot be written to that disk space because other drives do not have corresponding disk space available.

[Figure 4-2](#) shows an example of storage allocation in a RAID 5 array. The data is striped, with parity, across the three drives until the smallest drive is full. The remaining storage space in the other hard drives cannot be used because not all of the drives have disk space for redundant data.

 **NOTE:** Using hard disk drives of different sizes is not recommended.

Figure 4-2. Storage in a RAID 5 Array

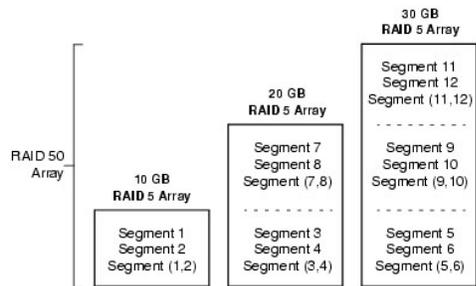


Storage in RAID 10 and RAID 50 Arrays

You can span RAID 1 and 5 arrays to create RAID 10 and RAID 50 arrays, respectively. For RAID levels 10 and 50, you can have some arrays with more storage space than others. After the storage space in the smaller arrays is full, you can use the additional space in larger arrays to store data.

[Figure 4-3](#) shows the example of a RAID 50 span with three RAID 5 arrays of different sizes. (Each array can have from three to 14 hard disks.) Data is striped across the three RAID 5 arrays until the smallest array is full. The data is striped across the remaining two RAID 5 arrays until the smaller of the two arrays is full. Finally, data is stored in the additional space in the largest array.

Figure 4-3. Storage in a RAID 50 Array



Performance Considerations

The system performance improves as the number of spans increases. As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 array.

Clearing Physical Drives

You can clear the data from SCSI drives using the configuration utilities. To clear a drive, perform the following steps:

1. Select **Management Menu**—> **Objects**—> **Physical Drives** in the BIOS Configuration Utility.

A device selection window displays the devices connected to the current controller.

2. Press the arrow keys to select the physical drive to be cleared and press <Enter>.
3. Select **Clear**.
4. When clearing completes, press any key to display the previous menu.

🔔 **NOTICE:** Do not terminate the clearing process, as it makes the drive unusable. You would have to clear the drive again before you could use it.

Displaying Media Errors

Check the **View Drive Information** screen for the drive to be formatted. Perform the following steps to display this screen which contains the media errors:

1. Select **Objects**—> **Physical Drives** from the **Management Menu**.
2. Select a device.
3. Press <F2>.

The error count displays at the bottom of the properties screen as they occur. If you feel that the number of errors is excessive, you should probably clear the hard drive. You do not have to select **Clear** to erase existing information on your SCSI disks, such as a DOS partition. That information is erased when you initialize logical drives.

Designating Drives as Hot Spares

Hot spares are physical drives that are powered up along with the RAID drives and usually stay in a standby state. If a hard drive used in a RAID logical drive fails, a hot spare will automatically take its place and the data on the failed drive is reconstructed on the hot spare. Hot spares can be used for RAID levels 1, 5, 10, and 50. Each controller supports up to eight hot spares.

🔔 **NOTE:** In the BIOS Configuration Utility, only global hot spares can be assigned. Dedicated hot spares cannot be assigned.

The methods for designating physical drives as hot spares are:

1. Pressing <F4> while creating arrays in **Easy**, **New** or **View/Add Configuration** mode.

- 1 Using the **Objects—> Physical Drive** menu.

<F4> Key

When you select any configuration option, a list of all physical devices connected to the current controller appears. Perform the following steps to designate a drive as a hot spare:

1. On the **Management Menu** select **Configure**, then a configuration option.
2. Press the arrow keys to highlight a hard drive that displays as **READY**.
3. Press <F4> to designate the drive as a hot spare.
4. Click **YES** to make the hot spare.

The drive displays as **HOTSP**.

5. Save the configuration.

Objects Menu

1. On the **Management Menu** select **Objects—> Physical Drive**.

A physical drive selection screen appears.

2. Select a hard drive in the **READY** state and press <Enter> to display the action menu for the drive.
3. Press the arrow keys to select **Make HotSpare** and press <Enter>.

The selected drive displays as **HOTSP**.

Rebuilding Failed Hard Drives

If a hard drive fails in an array that is configured as a RAID 1, 5, 10, or 50 logical drive, you can recover the lost data by rebuilding the drive with another drive or drives. You can manually rebuild one drive or a group of drives using the manual rebuild procedures in this section.

If a system is rebooted during a rebuild operation, it is possible for the rebuild to restart from 0 percent.

 **NOTE:** An array may take longer to rebuild when under high stress; for example, when there is one rebuild I/O operation to five host I/O operations.

 **NOTE:** In a clustering environment, if a node fails during a rebuild, the rebuild is re-started by another node. The rebuild on the second node starts at zero percent.

Rebuild Types

[Table 4-13](#) describes automatic and manual rebuilds.

Table 4-13. Rebuild Types

Type	Description
Automatic Rebuild	If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed disks. Select Objects—> Physical Drive to display the list of physical drives while a rebuild is in progress. The hot spare drive changes to REBLD A[array number]-[drive number] , indicating the hard drive is being replaced by the hot spare. For example, REBLD A01-02 indicates that the data is being rebuilt on hard drive 2 in array 1.
Manual Rebuild	Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed drives are available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

 **NOTE:** If a rebuild to a hot spare fails for any reason, the hot spare drive will be marked as "failed".

Use the following procedures to rebuild one failed drive manually in an individual mode or multiple drives in a batch mode.

Manual Rebuild – Rebuilding an Individual Drive

1. Select **Objects**→**Physical Drive** from the **Management Menu**.

A device selection window displays the devices connected to the current controller.

2. Designate an available drive as a hot spare before the rebuild starts.

See [Designating Drives as Hot Spares](#) for instructions on designating a hot spare.

3. Press the arrow keys to select the failed physical drive you want to rebuild, then press <Enter>.
4. Select **Rebuild** from the action menu and respond to the confirmation prompt.

Rebuilding can take some time, depending on the drive capacity.

5. When the rebuild is complete, press any key to display the previous menu.

Manual Rebuild – Batch Mode

1. Select **Rebuild** from the **Management Menu**.

A device selection window displays the devices connected to the current controller. The failed drives display as **FAIL**.

2. Press the arrow keys to highlight any failed drives to be rebuilt.
3. Press the spacebar to select the desired physical drives for rebuild.
4. After you select the physical drives, press <F10> and select **Yes** at the prompt.

The selected drives change to **REBLD**. Rebuilding can take some time, depending on the number of drives selected and the drive capacities.

5. When the rebuild is complete, press any key to continue.
6. Press <Esc> to display the **Management Menu**.

Checking Data Consistency

Select the Check Consistency option in the configuration utility to verify the redundancy data in logical drives that use RAID levels 1, 5, 10, and 50. (RAID 0 does not provide data redundancy.) The parameters of the existing logical drives appear and discrepancies are automatically corrected when the data is correct. However, if the failure is a read error on a data drive, the bad data block is reassigned and the data is re-generated.

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

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

Perform the following steps to run **Check Consistency**:

1. Access the **Management Menu**.
2. Select **Check Consistency**.
3. Press the arrow keys to highlight the desired logical drives.
4. Press the spacebar to select or deselect a drive for consistency checking.
5. Press <F2> to select or deselect all the logical drives.
6. Press <F10> to begin the consistency check.

A progress graph for each selected logical drive displays.

7. When the check is finished, press any key to clear the progress display.
8. Press <Esc> to display the **Management Menu**.

(To check an individual drive, select **Objects—> Logical Drives** from the **Management Menu**, the desired logical drive(s), then **Check Consistency** on the action menu.)

 **NOTE:** Stay at the **Check Consistency** menu until the check is complete.

Reconstructing Logical Drives: RAID Level Migration and Online Capacity Expansion

A reconstruction occurs when you change the RAID level of an array or add a physical drive to an existing array. RAID level migration changes the array from one RAID level to another. Online capacity expansion is the addition of hard disk drives to increase storage capacity. You can perform a reconstruction while the system continues to run, without having to reboot. This avoids downtime and keeps data available to users.

 **NOTE:** After you start the reconstruct process, you must wait until it is complete. Do not reboot, cancel, or exit until the reconstruction is complete.

 **NOTE:** When performing a RAID level migration or an online capacity expansion, a fictional disk may appear in the Windows Disk Management, Dell OpenManage Array Manager, or Dell OpenManage Storage Services application, if the system is rebooted before the process is finished. This disk can be ignored and will disappear once the RAID level migration or online capacity expansion is complete.

Performing a RAID level migration on a clustered system will change the system to non-clustered mode, causing a cluster mismatch error if the system is rebooted.

 **NOTE:** An automatic drive rebuild will not start if you replace a drive during a RAID level migration or an online capacity expansion. The rebuild must be started manually after the expansion or migration procedure is complete.

Perform the following steps to reconstruct a drive:

1. Move the arrow key to highlight **Reconstruct** on the **Management Menu**.
2. Press <Enter>.

A window entitled "**Reconstructables**" displays. This contains the logical drives that can be reconstructed. You can press <F2> to view logical drive information or <Enter> to select the reconstruct option.

3. Press <Enter>.

The next reconstruction window displays. The options on this window are <spacebar> to select or deselect a drive, <Enter> to open the reconstruct menu, and <F3> to display logical drive information.

4. Press <Enter> to open the reconstruct menu.

The menu items are RAID level, stripe size, and reconstruct.

5. To change the RAID level, select **RAID** with the arrow key, press <Enter> and select a RAID level from the list that displays.
6. Select **Reconstruct** and press <Enter> to reconstruct the logical drive.

You are prompted to start the reconstruction. A progress bar for the reconstruction displays.

Drive Roaming

Drive roaming occurs when the hard drives are changed to different target IDs or channels on the same controller. When the drives are placed on different channels, the controller detects the RAID configuration from the configuration data on the drives.

 **NOTE:** In a clustering environment, drive roaming is supported within the same channel only.

Configuration data is saved in both non-volatile random access memory (NVRAM) on the RAID controller and on the hard drives attached to the controller. This maintains the integrity of the data on each drive, even if the drives have changed their target ID.

 **NOTE:** If you move a drive that is currently being rebuilt, the rebuild operation will restart, not resume.

Perform the following steps to use drive roaming:

1. Turn off all power to the server and all hard drives, enclosures, and system components, then disconnect power cords from the system.
2. Open the host system by following the instructions in the host system technical documentation.
3. Move the drives to different positions on the backplane to change the SCSI ID.
4. Determine the SCSI ID and SCSI termination requirements.

 **NOTE:** The default for SCSI termination is onboard SCSI termination enabled.

5. Perform a safety check.
 1. Make sure the drives are inserted properly.
 1. Close the cabinet of the host system.
 1. Turn power on after completing the safety check.
6. Power on the system.

The controller then detects the RAID configuration from the configuration data on the drives.

Drive Migration

Drive migration is the transfer of a set of hard drives in an existing configuration from one controller to another. The drives must remain on the same channel and be reinstalled in the same order as in the original configuration. The controller to which you migrate the drives cannot have an existing configuration.

 **NOTE:** Only complete configurations can be migrated; individual virtual disks cannot be migrated.

 **NOTE:** Drive roaming and drive migration cannot be supported at the same time.

Perform the following steps to migrate drives:

1. Make sure that you clear the configuration on the system to which you migrate the drives, to prevent a configuration data mismatch between the hard drives and the NVRAM.

 **NOTE:** When you perform a drive migration, move only the disks that make up the logical drive (not all the physical disks in an array), so you will not see an NVRAM mismatch error (providing a configuration is on the destination controller). The NVRAM mismatch error appears only if you move all of the physical drives to the other controller.

2. Turn off all power to the server and all hard drives, enclosures, and system components, then disconnect power cords from the systems.
3. Open the host systems by following the instructions in the host system technical documentation.
4. Remove the unshielded, twisted pair, SCSI ribbon cable connectors from the internal drives or the shielded cables from the external drives you want to migrate.
 1. Make sure pin 1 on the cable matches pin 1 on the connector.
 1. Make sure that the SCSI cables conform to all SCSI specifications.
5. Remove the hard drives from the first system and insert them into drive bays on the second.
6. Connect the SCSI cables to the hard drives in the second system.
7. Determine the SCSI ID and SCSI termination requirements.

 **NOTE:** The default for SCSI termination is onboard SCSI termination enabled.

8. Perform a safety check.
 1. Make sure all cables are properly attached.
 1. Make sure the RAID controller is properly installed.

- 1 Close the cabinet of the host system.
 - 1 Turn power on after completing the safety check.
9. Power on the system.

The controller then detects the RAID configuration from the configuration data on the drives.

Deleting Logical Drives

This RAID controller supports the ability to delete any unwanted logical drives and use that space for a new logical drive. You can have an array with multiple logical drives and delete a logical drive without deleting the whole array.

After you delete a logical drive, you can create a new one. You can use the configuration utilities to create the next logical drive from a free space ('hole'), and from the newly created arrays. The configuration utility provides a list of configurable arrays where there is a space to configure. In the BIOS Configuration Utility, you must create a logical drive in the hole before you create a logical drive using the rest of the disk.

 **NOTE:** Warning messages display about the effect of deleting an array. You must accept two warning statements before the array deletion is completed.

 **NOTICE:** The deletion of the logical drive can fail under certain conditions: During a rebuild, initialization or check consistency of a logical drive.

To delete logical drives, perform the following steps in the BIOS Configuration Utility:

1. Select **Objects**—> **Logical Drive** from the **Management Menu**.

The logical drives display.

2. Use the arrow key to highlight the logical drive you want to delete.
3. Press to delete the logical drive.

This deletes the logical drive and makes the space it occupied available for you to make another logical drive.

Patrol Read

The Patrol Read function is designed as a preventive measure to detect hard drive errors before drive failure can threaten data integrity. Patrol Read can find and possibly resolve any potential problem with physical drives prior to host access. This can enhance overall system performance because error recovery during a normal I/O operation may not be necessary.

Patrol Read Behavior

The following is an overview of Patrol Read behavior:

1. Patrol Read runs on all disks on the adapter that are configured as part of an array including hot spares. Patrol Read will not run on unconfigured drives, which are drives that are not part of an array or that are in a ready state.
2. Patrol Read adjusts the amount of RAID controller resources dedicated to Patrol Read operations based on outstanding disk I/O. For example, if the server is busy processing I/O operation, then Patrol Read will use less resources to allow the I/O to take a higher priority.
3. Patrol Read operates on all configured physical drives on the controller and there is no method to deselect drives from the Patrol Read operations.
4. If the server reboots during a Patrol Read iteration, Patrol Read will restart from zero percent if in **Auto Mode**. In **Manual Mode**, Patrol Read does not restart upon a reboot. **Manual Mode** assumes you have selected a window of time dedicated to running Patrol Read and the server will be available during that time.

Configuration

You can use the BIOS Configuration Utility to configure Patrol Read. Dell OpenManage Array Manager and OpenManage System Storage Management cannot configure Patrol Read. Patrol Read can be started and stopped using MegaPR from within Window and Linux.

Blocked Operations

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

- 1 An unconfigured disk (the disk is in the READY state)
- 1 Disks that are members of a logical drive undergoing a reconstruction
- 1 Disks that are members of a logical drive that is currently owned by the peer adapter in a cluster configuration
- 1 Disks that are members of a logical drive undergoing a Background Initialization or Check Consistency

Patrol Read Scheduling Details

The following describes the scheduling details for Patrol Read:

1. The PERC controller default settings sets Patrol Read to **Auto** mode. The Patrol Read mode can be set to **Auto** or **Manual** mode in the BIOS Configuration Utility.
2. In **Auto** mode, Patrol Read runs continuously on the system and is scheduled to start a new Patrol Read within four hours after the last iteration is completed.
3. When **Patrol Read Mode** is changed from **Auto** to **Manual**, **Manual Halt**, or **Disabled**, the **Next execution will start at:** field will be set to **N/A**.

Configuring Patrol Read

Patrol Read can be set to **Manual** or **Automatic** mode. When in **Manual** mode, the BIOS Configuration Utility can start and stop a Patrol Read iteration. MegaPR can be used to start and stop a Patrol Read iteration from Linux or Windows.

The BIOS Configuration Utility has options to configure Patrol Read on the controller. Access the **Objects→ Adapter→ Patrol Read Options** menu. Press <Enter> to open the Patrol Read submenu, which displays the following items:

- 1 Patrol Read Mode
- 1 Patrol Read Status
- 1 Patrol Read Control

Patrol Read Mode

The current setting displays as **Manual/Auto/Disabled**. When you select this option, a window opens to display the following options, with the current setting highlighted:

1. Manual
2. Auto
3. Manual Halt
4. Disabled

You can change the setting by selecting a different value upon confirmation.

Patrol Read Status

When you select **Patrol Read Status** and press <Enter>, a window opens to display these options:

1. Number of Iterations Completed =

2. State = Active/Stopped
3. Next Execution will Start at

The current state is shown at the second option that allows you to display the percentage of completion by pressing <Enter> key if the Patrol Read state is Active. The first and the third options are read only.

Patrol Read Control

When you select this option, a window opens to display the following options:

1. Start
2. Stop

 **NOTE:** Start or Stop options are available in manual mode only.

Behavior Details

The following are behavior details for Patrol Read:

1. Setting Patrol Read in Manual mode does not start Patrol Read. It only sets the mode so that you can select start whenever you want to run Patrol Read. Once the mode is MANUAL, it stays in that mode until you change it.
2. Setting the mode to AUTOMATIC starts Patrol Read; when the Patrol Read operation is complete, it will set itself to run within four hours of the last iteration completion.

MegaPR Utility

MegaPR is a utility for managing and reporting the status of Patrol Read from the operating system. There are two versions of the utility: one for Windows 2000/2003, and one for Linux (RHEL 2.1, 3, and 4).

Available options are (help for individual options is available by typing `cmd -[option] ?`):

- l -dispPR: Display Patrol Read status.
- l -startPR: Start Patrol Read.
- l -stopPR: Stop Patrol Read.

[Back to Contents Page](#)

[Back to Contents Page](#)

Driver Installation

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [Obtaining Drivers](#)
- [Using the Dell OpenManage Installation and Server Management or Server Assistant CD to Install an Operating System](#)
- [Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette](#)
- [Installing a Windows 2000 or 2003 Driver for a New RAID Controller](#)
- [Updating an Existing Windows 2000 or 2003 Driver](#)
- [Installing the Linux RedHat Driver](#)
- [Installing the Novell NetWare Driver](#)
- [Modifying the PCI Slot Numbers for the Controllers](#)

The Dell™ PowerEdge Expandable RAID Controller (PERC) 4/Di/Si and 4e/Di/Si controllers require software drivers to operate with the Microsoft® Windows®, RedHat® Linux®, and Novell NetWare® operating systems.

The drivers support:

- 1 40 logical drives per RAID controller
- 1 The ability to detect newly configured logical drives in Disk Administrator without rebooting the system (applicable only to Windows operating systems)
- 1 The ability to delete the last logical drive created using the configuration utilities (See the RAID controller's user's guide for more information.)
- 1 The ability to use the remaining capacity of an array using Dell OpenManage Array Manager or Dell OpenManage Storage Management (if provided).

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

- 1 Microsoft Windows 2000/2003 Server
- 1 Red Hat Linux
- 1 Novell Netware

There are three methods for installing the driver:

- 1 During operating system installation. Use this method if you are performing a new installation of the operating system and want to include the drivers.
- 1 After adding a new RAID controller. Use this method if the operating system is already installed, you have installed a RAID controller, and you want to add the device drivers.
- 1 Updating existing drivers. Use this method if the operating system and RAID controller are already installed, and you want to update to the latest drivers.

Obtaining Drivers

 **NOTE:** See the readme file included with the driver for any updated information.

A driver diskette can be created for each supported operating system from the Dell OpenManage Installation and Server Management or Server Assistant. However, to make sure you have the latest version of drivers, download the updated drivers from the Dell Support web site at: <http://support.dell.com>.

Using the Dell OpenManage Installation and Server Management or Server Assistant CD to Install an Operating System

The *Dell Installation and Server Management* or *Dell Server Assistant* CD is a bootable, stand-alone CD-ROM that provides the tools required to setup and configure new Dell PowerEdge system components and software. It carries the latest available drivers that have been optimized for use on Dell PowerEdge Servers.

The *Dell Installation and Server Management* or *Dell Server Assistant* CD provides significant enhancements that streamline the installation of the Operating System on the PowerEdge Server. The *Dell Installation and Server Management* or *Dell Server Assistant* CD ships with every Dell PowerEdge Server. This proven set of tools and documentation greatly enhances the customer's out-of-box experience by providing an easy to follow step-by-step setup and operating system installation process.

Perform the following steps to install the driver while you are installing the operating system with the *Dell Installation and Server Management* or *Dell Server Assistant* CD.

1. Power the system down.
2. Power the system on.
3. During boot up, the PERC BIOS banner should display. If it does not, power the system down and refer to [Troubleshooting](#).
4. Configure the logical drives. For more information on setting up logical drives, refer to [RAID Configuration and Management](#).

 **NOTE:** If this controller is not your primary controller, you can skip to Step 6 and configure the logical drives using Dell OpenManage Array Manager or Dell OpenManage Storage Management (if provided) once the operating system is installed.

5. Insert the *Dell Installation and Server Management* or *Dell Server Assistant* CD in the CD drive and restart the server.
6. Select the language that you want to use when prompted.
7. Read and accept the software license agreement to continue.
8. Select Click here for Server Setup on the Systems Management main page.
9. Follow the instructions on the screen to complete setting up the operating system.
10. Systems Management detects the devices on your system and then automatically installs drivers for all of those devices, including your RAID controller.
11. When prompted, insert the operating system CD and follow the instructions on the screen to complete the installation. Refer to the operating system documentation for more information on completing the operating system installation.

Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette

Creating a Driver Diskette

A driver diskette can be created through one of the following two methods:

1. Obtain the driver from the *Dell OpenManage Systems Management* CD or *Support* CD
1. Obtain the latest drivers from Dell Support located at: <http://support.dell.com>.

To create a driver diskette using the Dell OpenManage Systems Management CD or Support CD, perform the following steps:

1. Insert the *Dell OpenManage Systems Management* CD or *Support* CD, into a running systems CD drive, and insert a diskette into the diskette drive.
2. After the CD autoruns, click **Copy Drivers**.
3. Select a server from the Select Server drop-down menu, then select the operating system under Select Drivers/Utilities Set.
4. Click **Continue**.
5. On the Utilities and Drivers page, scroll to the box for the operating system for the server and click the driver for your type of RAID controller.
6. Follow the instructions on the screen and unzip the file to the diskette.

To create a driver diskette using the Dell support site, perform the following steps:

1. Browse to the download section for the server at: <http://support.dell.com>.
2. Locate and download the latest RAID driver to the system. The driver should be labeled as packaged for a diskette on the support site.
3. Follow the instructions on the support site for extracting the driver to the diskette.

Installing the Driver during Operating System Installation

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

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

3. Press the <S> key.

The system prompts for the driver diskette to be inserted.

4. Insert the driver diskette in the floppy drive and press the <Enter> key.

A list of PERC controllers appears.

5. Select the right driver for the installed controller and press the <Enter> key to load the driver.

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

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

Installing a Windows 2000 or 2003 Driver for a New RAID Controller

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

1. Power down the system.
2. Install the new RAID controller in the system.

Refer to [RAID Configuration and Management](#) for detailed instructions on installing and cabling the RAID controller in the system.

3. Power on the system.

The Windows operating system should detect the new controller and display a message to inform the user.

4. The Found New Hardware Wizard screen pops up and displays the detected hardware device.
 5. Click **Next**.
 6. On the **Locate device driver** screen, select **Search for a suitable driver for my device** and click Next.
 7. Insert the appropriate driver diskette and select **Floppy disk drives** on the **Locate Driver Files** screen.
 8. Click **Next**.
 9. The wizard detects and installs the appropriate device drivers for the new RAID controller.
 10. Click **Finish** to complete the installation.
 11. Reboot the server.
-

Updating an Existing Windows 2000 or 2003 Driver

Perform the following steps to update the windows driver for the RAID controller already installed on your system.

 **NOTE:** It is important that you idle your system before you update the driver.

1. Press **Start** → **Settings** → **Control Panel** → **System**.

The **System Properties** screen displays.

 **NOTE:** In Windows 2003, press **Start** → **Control Panel** → **System**.

2. Click on the **Hardware** tab.
3. Click the **Device Manager** and the **Device Manager** screen displays.
4. Click **SCSI** and **RAID** Controllers.
5. Double-click the RAID controller for which you want to update the driver.
6. Click the **Driver** tab and click on **Update Driver**.

The screen for the **Upgrade Device Driver** Wizard displays.

7. Insert the appropriate driver diskette.
8. Click **Next**.
9. Follow the steps in the Wizard to search the diskette for the driver.
10. Select the INF file from the diskette.

 **NOTE:** In Windows 2003, select the name of the driver, not the INF file.

11. Click **Next** and continue the installation steps in the Wizard.
12. Click **Finish** to exit the wizard and reboot the system for the changes to take place.

Installing the Linux RedHat Driver

Use the procedures in this section to install the RedHat Linux driver for RedHat Linux 8.1, 9.0, AS 2.1, 3.0, and ES 2.1, 3.0. The driver is updated frequently. To make sure you have the current version of the driver, you can download the updated RedHat Linux driver from Dell Support at support.dell.com.

 **NOTE:** On a Linux 8.0 system, when you run Cerc Manager (v. 5.23) from a Gnome-terminal in XWindows, the <F10> key cannot be used to create a logical drive. Instead, you can use the alternate keys <Shift><0>. (This is not an issue if Xterm is used to call cercmgr). The following is a list of alternate keys you can use in case of problems with keys <F1> through <F6>, and <F10>:

- n <Shift><1> for <F1>
- n <Shift><2> for <F2>
- n <Shift><3> for <F3>
- n <Shift><4> for <F4>
- n <Shift><5> for <F5>
- n <Shift><6> for <F6>
- n <Shift><0> for <F10>

To install a RedHat Linux driver more recent than the one on the RedHat CD, you must use a driver diskette when you are installing the operating system. See [Installing the Driver](#) for information on this procedure. You must download the files before you begin the operating system installation.

For more detailed installation instructions for RedHat Linux 9.0 or later, see the operating system installation guide on the Dell Support site at: support.dell.com.

Creating a Driver Diskette

Before beginning the installation, download the driver appropriate for your version of RedHat Linux from support.dell.com to your temporary directory. This file includes two RPMs and five driver disk files. From a RedHat Linux system, enter the following commands to separate the individual driver files from the tar archive file:

```
mount /dev/fd0 /mnt/floppy
```

```
tar xvzf -C /mnt/floppy /tmp/filename.tar.gz
```

 **NOTE:** You can also create a driver diskette using the *Dell OpenManage Systems Management CD* or *Server Support CD*. See [Creating a Driver Diskette](#) in the [Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette](#) section for more information.

Installing the Driver

Perform the following steps to install RedHat Linux 9.0 or later and the appropriate RAID drivers.

1. Boot normally from the RedHat Linux installation CD.
2. At the command prompt, type:

```
expert noprobe dd
```

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

See [Creating a Driver Diskette](#) for information about creating a driver diskette.

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

Installing the Driver Using an Update RPM

The following procedures explain how to install RedHat Linux 9.0 or later and the appropriate RAID driver using an update RPM with or without DKMS support.

Installing the RPM Package without DKMS Support

Perform the following steps to install the RPM package without DKMS support:

1. Download the driver rpm package from support.dell.com.
2. Copy the driver rpm package to the proper location.
3. Install the driver rpm package:

```
rpm -Uvh <driver_package_file_name>
```

4. Reboot the system to load the new driver.

Installing the RPM Package with DKMS Support

Perform the following steps to install the RPM package without DKMS support:

1. Decompress the zipped file of the DKMS-enabled driver package.
2. In the directory containing the decompressed file, type the following shell command:

```
sh install.sh
```

3. Reboot the system to load the new driver.
4. Create driver diskette image using DKMS.

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

The following files are needed before creating the DUD.

 **NOTE:** The megaraid2 driver package installs these files. You do not need to do anything at this point.

1. There is a directory `/usr/src/megaraid2-<driver_version>`, which contains the driver source code, `dkms.conf` and spec file for the driver.
2. In this directory, there is a subdirectory called `redhat_driver_disk` which contains the files needed to create the DUD. The files needed are: `disk_info`, `modinfo`, `modules.dep`, and `pcitable`.
3. To create the DUD image for pre-RedHat4 distribution, the kernel source package has to be installed to compile the driver. For RedHat4 distribution, the kernel source is not needed.

DUD Creation Procedure

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

1. Install the DKMS-enabled megaraid2 driver rpm package on a RedHat system.
2. Type the following command in any directory:

```
dkms mkdriverdisk -d redhat -m megaraid2 -v <driver version> -k <kernel version>
```

This starts the process to create the megaraid2 DUD image.

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

```
dkms mkdriverdisk -d redhat -m megaraid2 -v <driver version> -k <kernel_version_1>, <kernel_version_2>, ...
```

4. After the DUD image has been built, you can find it in the DKMS tree for megaraid2 driver.

 **NOTE:** Currently the DKMS package supports creating the DUD in RedHat distribution only. You can create a DUD on RedHat only.

Installing the Novell NetWare Driver

You can use the following methods to install the Novell NetWare® drivers:

- 1 During operating system installation

Use this method if you are performing a new installation of Novell NetWare using Dell Systems Management and want to include the drivers. See [Installing the Driver during Operating System Installation](#) in [Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette](#) for more information.

 **NOTE:** For information about installing drivers if you use the NetWare CD to install your operating system, see your Novell documentation.

- 1 After adding a new RAID controller

Use this method if Novell NetWare is already installed and you want to add the device drivers after installing the RAID controller.

- 1 Performing a Standard Mode Installation of NetWare 5.1SBE, 6.0, and 6.5

With standard mode installation, you accept the defaults for the components to be installed.

- 1 Updating existing drivers

Use this method if Novell NetWare and the RAID controller are already installed, and you want to update to the latest drivers for the controller.

Installing a NetWare Driver for a New Controller

Perform the following steps to add a NetWare 5.1, 6.0, 6.5, or later driver to an existing installation.

1. At the root prompt, perform the following steps:
 - a. For NetWare 5.1 and 6.0, type:

```
nwconfig
```

and press <Enter>.

The **Installation Options** screen displays.

- b. For NetWare 6.5, type:

```
hdetect
```

and press **Continue** on the first menu to go to the storage drivers, then follow the instructions for updating the driver. For NetWare 6.5, you can press <F3> to auto detect drivers.

2. Select **Configure Disk** and **Storage Device Options**, then press <Enter>.
3. Select one of the options that display:
 - i Discover and load an additional driver.

If you select the option **Discover and load an additional driver**, the system detects the extra unit. Perform [step 4](#) to complete the procedure.

4. At the prompt to select a driver from the list, press <Insert> to insert the driver, which completes the procedure.

If you select the option **Select an additional driver**, perform steps 5 - 8.

5. After you select **Select an additional driver**, the **Select a Driver** screen displays.
6. Press <Insert> and read the instructions that display.
7. Put the driver diskette in the diskette drive and press <Enter>.
8. The system then detects a driver and installs it.

Modifying the PCI Slot Numbers for the Controllers

Perform the following steps to modify the PCI slot numbers for the controller:

1. At the command prompt, type:

```
C:\NWSERVER>
```

and press <Enter>.

2. Type

```
server -nss
```

(Not load Storage Service /modules.NLM)

3. At : Prompt (System Console), type:

```
load pedge3.ham
```

and press <Enter>.

The following supported slot options display:

- 1 No Selection
- 1 PCI Slot_2.1 (HIN 202)
- 1 PCI EMBEDDED (HIN 10017)

 **NOTE:** Write down the number after "HIN". In step 3, the number is 10017.

4. Under choice, type:

0

This is for no selection.

5. At the command prompt (System Console), type

```
Edit Startup.ncf
```

A list of CDM drivers displays.

6. Select `LOAD PEDGE3.HAM SLOT=XXXX`.
7. Before you exit the list of CDM drivers, press <Alt> <V> to save the update.
8. Press <Alt> <X> to exit to C:\NWSERVER.
9. At the C:\NWSERVER prompt, type the following for the operating system to boot:

```
server
```

The operating system boots.

Performing a Standard Mode Installation of NetWare 5.1SBE, 6.0, and 6.5

Standard mode means that you accept the defaults for the components to be installed. Perform the following steps for a standard mode installation on NetWare 5.1SBE, 6.0, and 6.5:

1. At **Server Settings**, select **Continue** and press <Enter> to accept the default.
2. At **Regional Settings**, select **Continue** and press <Enter> to accept the default.
3. At **Mouse type** and **Video mode** select **Continue** and press <Enter> to accept the default.

The system will take several minutes to load files. It will find the device drivers that support the adapter.

4. Insert the driver diskette in the floppy (A:/) drive.
5. For device types and driver names, select **Modify** and press <Enter>.
6. Highlight **Storage Adapters** and press <Enter>.
7. At the option **Add, Edit or Delete Storage Drivers**, press <Insert> to add a driver.
8. At the option **Select a Driver for each Storage Adapter**, press <Insert> to add an unlisted driver.

The system scans the path for the A:/ drive. The driver diskette is already in the A:/ drive. The option **Return to Driver Summary** displays.

9. Select **Return to Driver Summary** and press <Enter>.

10. Select **Continue** and press <Enter>.

 **NOTE:** You must load a driver for each controller. For example, if you have four adapters, the driver is listed four times.

Updating an Existing Driver for NetWare 5.1 or 6.0

Perform the following steps to update an existing driver for NetWare 5.1 or 6.0:

1. Create a driver diskette.

See [Creating a Driver Diskette](#) in the [Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette](#) section for information. (The procedure for creating a driver diskette is the same for all operating systems.)

2. Once the NetWare server is up, type the following:

```
nwconfig
```

3. Press <Enter> to access the NetWare Configuration Utility.
4. On the **Configuration Options** screen, select **Driver Options** and press <Enter>.
5. Under the **Driver Options**, select **Configure Disk and Storage Options**, then press <Enter>.
6. Under the **Additional Driver Actions** menu, press the down arrow key to select the **Additional Driver** option, then press <Enter>.
7. Press <Insert> to install an unlisted driver.
8. Press <Insert> again if using a diskette; otherwise, press <F3> to specify a different location.
9. Insert the driver diskette into the diskette drive and press <Enter>.

The file **pedge3.ham** displays under the option **Select a Driver to Install**.

10. Highlight **pedge3.ham** and press <Enter>.
11. Select **Yes** to copy **pedge3.ham** files to C:\NWSERVER.
12. Select **No** to save the existing file messages to C:\NWSERVER.
13. Under **pedge3 Parameters**, perform the following steps to provide the slot number.
14. Press <Alt><Esc> to access **System Console**.
15. On the **System Console**, type:

```
load pedge3
```

16. Press <Enter>.

The following supported slot options display:

```
1 No Selection
1 PCI Slot_2.1 (HIN 203)
```

17. Write down the number after "HIN".

In the example in [step 16](#), it is 203.

18. Under **Choice**, type:

```
0
```

for the option **No Selection**.

19. Unload **pedge3.ham**.
20. Press <Alt><Esc> until you exit the System Console and return to the **pedge3 Parameters** screen in the NetWare Configuration Utility.

21. Under **Slot Number**, enter the slot number you obtained from System Console and press <Enter>.
22. Press <F10> to save the **pedge3** parameters.
23. Under **Driver pedge3 Parameters Actions**, select **Save Parameters and Load Driver**, and press <Enter>.
24. Select **No** when asked to load additional drivers.

pedge3 will be listed on the **Selected Disk Driver** screen.

25. Exit the NetWare Installation Utility.
26. From server console, type:

```
reset server
```

to restart the server for the changes to take effect.

Updating an Existing Driver for NetWare 6.5

Perform the following steps to update an existing driver for NetWare 6.5:

1. Create a driver diskette.

See [Creating a Driver Diskette](#) in the [Installing Windows 2000 or 2003 Using the Microsoft Operating System CD with Driver Diskette](#) section for information.

2. Once NetWare begins to boot, the following message displays: **Press ESC to abort OS boot.**
3. Press <Esc>.
4. At the command prompt, type:

```
C:\NWSERVER>
```

and press <Enter>.

5. At C:\NWSERVER>, insert a driver diskette into the floppy drive.
6. Type:

```
cd A:\
```

and press <Enter>.

7. At the A:\ prompt, type:

```
copy A:\*. * C:\NWSERVER\DRIVERS
```

and wait for the copy action to complete.

8. Change the directory and type:

```
cd C:
```

The path displays as "C:\NWSERVER".

9. At C:\NWSERVER prompt, type:

```
server
```

The operating system starts to boot.

10. To verify the driver version, at System Console (1) type:

```
modules Pedge3*
```

The driver version displays.

[Back to Contents Page](#)

[Back to Contents Page](#)

Troubleshooting

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

- [Logical Drive Degraded](#)
- [System CMOS Boot Order](#)
- [General Problems](#)
- [Hard Disk Drive Related Issues](#)
- [Drive Failures and Rebuilds](#)
- [SMART Error](#)
- [BIOS Error Messages](#)

To get help with problems with your RAID controller, you can contact your Dell Service Representative or access the Dell Support web site at support.dell.com.

Logical Drive Degraded

A logical drive is in a degraded condition when one hard drive in its span has failed or is offline. For example, when a RAID 10 logical drive consisting of two spans of two drives each can sustain a drive failure in each span and be a degraded logical drive. The RAID controller has the fault tolerance to undergo a single failure in each span without compromising data integrity or processing capability.

The RAID controller provides this support through redundant arrays in RAID levels 1, 5, 10 and 50. The system can still work properly even with a single disk failure in an array, though performance can be degraded to some extent.

To recover from a degraded logical drive, rebuild the failed drive in each array. Upon successful completion of the rebuild process, the logical drive state changes from degraded to optimal. For the rebuild procedure, see [Rebuilding Failed Hard Drives](#) in [RAID Configuration and Management](#).

System CMOS Boot Order

If you intend to boot to the controller, ensure it is set appropriately in the system's CMOS boot order. Refer to the system documentation for your individual system.

 **NOTE:** Only the first eight logical drives can be used as bootable devices.

General Problems

[Table 6-1](#) describes general problems you might encounter, along with suggested solutions.

Table 6-1. General Problems

Problem	Suggested Solution
The device displays in Device Manager but has a yellow bang (exclamation point).	Reinstall the driver. See the driver installation procedures in Driver Installation .
Windows driver does not appear in Device Manager.	Power off the system and reset the card.
"No Hard Drives Found" message appears during a CD-ROM installation of Windows 2000 or Windows 2003 because of the following causes: <ol style="list-style-type: none">1. The drive is not native in the operating system.2. The logical drives are not configured properly.3. The controller BIOS is disabled.	The corresponding solutions to the three causes of the message are: <ol style="list-style-type: none">1. Press <F6> to install the RAID Device Driver during installation.2. Enter the BIOS Configuration Utility to configure the logical drives. See RAID Configuration and Management for procedures to configure the logical drives.3. Enter the BIOS Configuration Utility to enable the BIOS. See RAID Configuration and Management for procedures to configure the logical drives.

<p>The BIOS Configuration Utility does not detect a replaced physical drive in a RAID 1 array and offer the option to start a rebuild.</p> <p>After the drive is replaced, the utility shows all drives online and all logical drives reporting optimal state. It does not allow rebuilding because no failed drives are found.</p>	<p>This occurs if you replace the drive with a drive that contains data. If the new drive is blank, this problem does not occur.</p> <p>Perform the following steps to solve this problem:</p> <ol style="list-style-type: none"> 1 Access the BIOS Configuration Utility and select Objects—> Physical Drive to display the list of physical drives. 1 Use the arrow key to select the newly inserted drive, then press <Enter>. <p>The menu for that drive displays.</p> <ol style="list-style-type: none"> 1 Select Force Offline and press <Enter>. <p>This changes the physical drive from online to Failed.</p> <ol style="list-style-type: none"> 1 Select Rebuild and press <Enter>. <p>After rebuilding is complete, the problem is resolved and the operating system will boot.</p>
<p>The system takes a long time to boot during a RAID Level Migration or Check Consistency operation.</p>	<p>This is normal behavior during a RAID level migration or consistency check.</p>

Hard Disk Drive Related Issues

[Table 6-2](#) describes hard drive related problems you might encounter, along with suggested solutions.

Table 6-2. Hard Disk Drive Issues

Problem	Suggested Solution
<p>The system does not boot from the RAID controller.</p>	<p>If the system does not boot from the controller, check the boot order in the BIOS.</p>
<p>One of the hard drives in the array fails often.</p>	<p>This could result from one or two problems.</p> <ol style="list-style-type: none"> 1 If the same drive fails: <ul style="list-style-type: none"> o Format the drive. o Check the enclosure or backplane for damage. o Check the SCSI cables. o Replace the hard drive. 1 Drives in the same slot keep failing: <ul style="list-style-type: none"> o Check the enclosure or backplane for damage. o Check the SCSI cables. o Replace the cable or backplane.
<p>Critical Array Status Error is reported during boot-up.</p>	<p>One or more of your logical drives is degraded. To recover from a degraded logical drive, rebuild the failed drive in each array. Upon successful completion of the rebuild process, the logical drive state changes from degraded to optimal. See Logical Drive Degraded in this section for more information. See Rebuilding Failed Hard Drives in RAID Configuration and Management for information about rebuilding failed drives.</p>
<p>FDISK reports much lower drive capacity in the logical drive.</p>	<p>Some versions of FDISK (such as DOS 6.2) do not support large disk drives. Use a version that supports large disk sizes or use a disk utility in your operating system to partition your disk.</p>
<p>Cannot rebuild a fault tolerant array.</p>	<p>This could result from any of the following:</p> <ol style="list-style-type: none"> 1 The replacement disk is too small or bad. Replace the failed disk with a good drive. 1 The enclosure or backplane could be damaged. Check the enclosure, or backplane. 1 The SCSI cables could be bad. Check the SCSI cables.
<p>Fatal errors or data corruption are reported when accessing arrays.</p>	<p>Contact Dell Technical Support.</p>

Drive Failures and Rebuilds

[Table 6-3](#) describes issues related to drive failures and rebuilds.

Table 6-3. Drive Failure and Rebuild Issues

Issue	Suggested Solution
-------	--------------------

Rebuilding a hard disk drive after a single drive failure	<p>If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed disks. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed drives are available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive. You can use the BIOS Configuration Utility or Dell OpenManage® Array Manager to perform a manual rebuild of an individual drive.</p> <p>Refer to Rebuilding Failed Hard Drives in RAID Configuration and Management for procedures for rebuilding a single hard disk drive.</p>
Rebuilding hard disk drives after a multi-drive failure	<p>Multiple drive errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data. It is possible to recover the logical drive from a multiple drive failure. Perform the following steps to recover the logical drive:</p> <ol style="list-style-type: none"> 1. Shut down the system, check cable connections, and reset hard drives. <p>Be sure to follow safety precautions to prevent electrostatic discharge.</p> <ol style="list-style-type: none"> 2. If the system logs are available, try to identify the order in which the drives failed in the multiple drive failure scenario. 3. Force the first drive online, then the second (if applicable), and continue till you reach the last disk. 4. Perform a rebuild on the last disk. <p>You can use the BIOS Configuration Utility or Dell OpenManage® Array Manager to perform a manual rebuild of multiple drives.</p> <p>See Rebuilding Failed Hard Drives in RAID Configuration and Management for procedures to rebuild a single hard disk drive.</p>
A drive is taking longer than expected to rebuild.	An array may take longer to rebuild when under high stress; for example, when there is one rebuild I/O operation for every five host I/O operations.
A node in a clustering environment fails during a rebuild.	In a clustering environment, if a node fails during a rebuild, the rebuild is re-started by another node. The rebuild on the second node starts at zero percent.

SMART Error

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

Table 6-4. SMART Error

Problem	Suggested Solution
A SMART error is detected in a fault-tolerant RAID array.	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. Force the hard disk drive offline. 2. Replace it with a new drive. 3. Perform a rebuild. <p>See Rebuilding Failed Hard Drives in RAID Configuration and Management for rebuild procedures.</p>
A SMART error is detected in non-fault-tolerant RAID array.	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. Back up your data. 2. Delete the logical drive. <p>See Deleting Logical Drives in RAID Configuration and Management for the procedure for deleting a logical drive.</p> <ol style="list-style-type: none"> 3. Replace the affected hard disk drive with a new drive. 4. Recreate the logical drive. <p>See Simple Array Setup or Advanced Array Setup in RAID Configuration and Management for procedures for creating logical drives.</p> <ol style="list-style-type: none"> 5. Restore the backup.

BIOS Error Messages

In PERC RAID controllers, the BIOS (option ROM) provides INT 13h functionality (disk I/O) for the logical drives connected to the controller, so that you can boot from or access the drives without the need of a driver. [Table 6-5](#) describes the error messages and warnings that display for the BIOS.

Table 6-5. BIOS Errors and Warnings

Message	Meaning
	This warning displays after you disable the option ROM in the configuration utility so that the BIOS will not hook Int13h and thus will not provide any I/O functionality to the logical drives.

BIOS Disabled. No Logical Drives Handled by BIOS	
Press <Ctrl><M> to Enable BIOS	When the BIOS is disabled, you are given the option to enable it by entering the configuration utility. You can change the setting to enabled in the configuration utility.
Configuration of NVRAM and drives mismatch Run View/Add Configuration option of Configuration Utility Press a key to enter Configuration Utility	If your boot-time BIOS options are set to Auto mode for BIOS configuration autoselection, the BIOS detects a mismatch of configuration data on the NVRAM and disks and this warning displays. You have to enter the configuration utility to resolve the mismatch before continuing. Perform the following steps to resolve the mismatch: 1. Press <Ctrl><M> to enter the BIOS Configuration Utility. 2. Select Configure—> View/Add Configuration from the Management Menu . The options Disk or NVRAM display. 3. Select either Disk to use the configuration data on the hard disk or NVRAM to use the configuration on the NVRAM. NOTICE: This message will display if any changes are made to the logical disk configuration in a clustered environment while one node is not in an up state. Accept the configuration from the disk.
Adapter at Baseport xxxx is not responding where xxxx is the baseport of the adapter	If the adapter does not respond for any reason but is detected by the BIOS, it displays this warning and continues. Shut down the system and try to reset the card. If this message still occurs, contact Dell Technical Support.
Insufficient Memory to Run BIOS. Press a Key to Continue	The BIOS needs some memory at POST to run properly. The BIOS allocates this memory either using PMM or another method. If the BIOS still cannot allocate the memory, it stops execution, displays this warning, then continues. This warning is very rare.
Insufficient Memory on the Adapter for the Current Configuration	If there is insufficient memory installed on the adapter, this warning displays and the system continues with another adapter. You should check to make sure the memory is properly installed and sufficient. Shut down the system and try to reset the card. If this message still occurs, contact Dell Technical Support.
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: 1 the adapter detects that the cache in the controller cache has not yet been written to the disk subsystem 1 the boot block detects an ECC error while performing its cache checking routine during initialization 1 the controller then discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed To resolve this problem, allow the battery to charge fully. If the problem persists, the battery or adapter DIMM might be faulty. In that case, contact Dell Technical Support.
x Logical Drive(s) Failed where x is the number of logical drives failed.	When the BIOS detects logical drives in the failed state, it displays this warning. You should check to determine why the logical drives failed and correct the problem. No action is taken by the BIOS.
x Logical Drives Degraded where x is the number of logical drives degraded.	When the BIOS detects logical drives in a degraded state, it displays this warning. You should try to make the logical drives optimal. No action is taken by the BIOS.
Following SCSI ID's are not responding Channel- ch1: id1, id2, Channel- ch2: id1, id2, where chx is channel number and id1 is first id that failed, id2 is second and so on.	When the BIOS determines that previously configured physical drives are not connected to the adapter, the BIOS displays this warning. You can connect the devices or take some other corrective action. The system continues to boot.
Adapter(s) Swap detected for Cluster/Non-Cluster mismatch	This warning displays when the BIOS detects a cluster/non-cluster mismatch in a cluster environment.
Warning: Battery voltage low	When the battery voltage is low, the BIOS displays this warning. You should check the battery.

Warning: Battery temperature high	When the battery temperature is high, the BIOS displays this warning. Your system is too hot. Check the air temperature and remove any obstructions to airflow. See messages below link.
Warning: Battery life low	Your RAID battery has a maximum number of charge and discharge cycles. When the BIOS displays this warning, the battery has reached the maximum number of cycles. Replace the battery.
<p>Following SCSI ID's have same data</p> <p>Channel- ch1: id1, id2,</p> <p>Channel- ch2: id1, id2,</p> <p>.</p> <p>where chx is channel number and id1 is first id that has same data, id2 is second and so on.</p>	This message displays when you perform drive roaming and the SCSI IDs have the same data.
<p>Error: Following SCSI Disk not found and No Empty Slot Available for mapping it</p> <p>No mapping done by firmware</p> <p>Channel- ch1: id1, id2,</p> <p>Channel- ch2: id1, id2,</p> <p>.</p> <p>where chx is channel number and id1 is first id that was not found, id2 is second and so on.</p>	This message displays when you perform drive roaming and no empty slot is available for the drive(s).

[Back to Contents Page](#)

[Back to Contents Page](#)

Glossary

Dell™ PowerEdge™ Expandable RAID Controller 4/Di/Si and 4e/Di/Si User's Guide

[A](#) • [B](#) • [C](#) • [D](#) • [E](#) • [G](#) • [H](#) • [I](#) • [L](#) • [M](#) • [N](#) • [O](#) • [P](#) • [R](#) • [S](#) • [T](#) • [W](#)

Adapter Swapping

When an adapter fails, a replacement can be inserted and connected to the existing set of drives. Dell supports adapter swapping only when all the attached disks are migrated to a new adapter that has a clear configuration.

Array

An array of disk drives combines the storage space on the disk drives into a single segment of storage space. The RAID controller can group disk drives on one or more SCSI channels into an array. A hot spare drive does not actively participate in an array.

Array Spanning

Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. The logical drives can span consecutively numbered arrays that each consist of the same number of disk drives. Array spanning promotes RAID level 1 to RAID level 10 and RAID level 5 to RAID level 50.

Asynchronous Operations

Operations that bear no relationship to each other in time and can overlap. The concept of asynchronous I/O operations is central to independent access arrays in throughput-intensive applications.

BIOS

(Basic Input/Output System) The part of the operating system in an IBM PC-compatible system that provides the lowest level interface to peripheral devices. The BIOS is stored in ROM in every IBM or compatible PC. BIOS also refers to the Basic Input/Output System of other "intelligent" devices, such as RAID controllers.

Cached I/O

Specifies that reads are buffered in cache memory, but does not override the other cache policies, such as read ahead or write.

Caching

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

Channel

An electrical path for the transfer of data and control information between a disk and a disk adapter. A channel can also be referred to as a "bus", such as a SCSI bus.

Clearing

In the BIOS Configuration Utility, the option used to delete data from physical drives.

Cold Swap

The replacement or exchange of a device in a system after powering down the system. In reference to disk subsystems, a cold swap requires that you turn the power off before replacing a defective hard drive.

Consistency Check

An examination of the disk system to determine if all conditions are valid for the specified configuration (such as parity).

Data Transfer Capacity

The amount of data per unit time moved through a channel. For disk I/O, bandwidth is expressed in megabytes per second (MB/s).

Degraded Drive

A disk drive that has become non-functional or has decreased in performance.

Direct I/O

Specifies that reads are not buffered in cache memory, but does not override the other cache policies, such as read ahead or write.

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 controlled by array management software. The array management software controls the disks and presents them to the array operating environment as a virtual disk.

Disk Duplexing

A variation on disk mirroring in which a second disk adapter or host adapter and redundant disk drives are present.

Disk Mirroring

Writing duplicate data to more than one (usually two) hard disks to protect against data loss in the event of device failure. Disk mirroring is a common feature of RAID systems.

Disk Spanning

The process of creating one logical drive composed of multiple arrays. Spanning is used to create complex RAID sets, such as RAID levels 10 and 50. Spanning utilizes striping to distribute data across all member disk drives.

Disk Striping

A type of disk array mapping. Consecutive stripes of data are mapped round-robin to consecutive array members. A striped array (RAID level 0) provides high I/O performance at low cost, but provides less data reliability than any member disk.

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 adapter, or the disks can attach directly to a system I/O bus adapter.

Double Buffering

A technique that achieves maximum data transfer bandwidth by constantly keeping two I/O requests outstanding for adjacent data. A software component begins a double-buffered I/O stream by issuing two requests in rapid sequence. Thereafter, each time an I/O request completes, another is immediately issued. If the disk subsystem can process requests fast enough, double buffering allows data to be transferred at the full-volume transfer rate.

Drive Roaming

Drive roaming occurs when the hard drives are changed to different target IDs or channels on the same controller. (A single-channel adapter can perform drive roaming.) When the drives are placed on different channels or target IDs, the controller detects the RAID configuration from the configuration information on the drives. Configuration data is saved in both NVRAM on the RAID controller and in the hard drives attached to the controller. This maintains the integrity of the data on each drive, even if the drives have changed their target ID.

Failed Drive

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

Fast SCSI

A variant on the SCSI-2 bus. It uses the same 8-bit bus as the original SCSI-1 but runs at up to 10MBs (double the speed of SCSI-1).

Firmware

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

Format

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

GB

A gigabyte: 1,000,000,000 (10 to the ninth power) bytes.

Host System

Any system to which disks are directly attached, (not attached remotely.) Mainframes, workstations, and personal computers can all be considered host systems.

Hot Spare

An idle, powered on, stand-by drive ready for immediate use in case of disk failure. It does not contain any user data. Up to eight disk drives can be assigned as hot spares for an adapter. A hot spare can be dedicated to a single redundant array or it can be part of the global hot-spare pool for all arrays controlled by the adapter.

When a disk fails, the controllers' firmware automatically replaces and rebuilds the data from the failed drive to the hot spare. Data can be rebuilt only from logical drives with redundancy (RAID levels 1, 5, 10, or 50; not RAID 0), and the hot spare must have sufficient capacity. The system administrator can replace the failed disk drive and designate the replacement disk drive as a new hot spare.

Hot Swap

The manual replacement of a failed drive while the disk subsystem is running (performing its normal functions).

Hot Swap Disk Drive

Hot swap drives allow a system administrator to replace a failed disk drive in a system without powering down the system and suspending services. The hot swap drive is pulled from its slot in the drive cage; all power and cabling connections are integrated into the drive enclosure backplane. The replacement hot swap drive can then slide into the slot. Hot swapping only works for RAID 1, 5, and 10 configurations.

Initialization

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

I/O Driver

A host system software component (usually part of the operating system) that controls the operation of peripheral adapters attached to the host system. I/O drivers communicate between applications and I/O devices and in some cases participate in data transfers.

Logical Drive

A complete or partial representation of a logical array. The storage space in a logical drive is spread across all the physical drives in the array or spanned arrays. Each RAID controller can be configured with up to forty logical drives in any combination of sizes. Configure at least one logical drive for each array. A logical drive can be in one of three states:

- 1 **Online:** All participating disk drives are online.
- 1 **Degraded:** (Also "Critical") a single drive in a redundant array (not RAID 0) is not online. Data loss can occur if a second disk drive fails.
- 1 **Offline:** Two or more drives in a redundant array (not RAID 0) or one or more drives in a RAID 0 array are not online.

I/O operations can be performed only with logical drives that are online or degraded.

Mapping

The relation between multiple data addressing schemes, especially conversions between member disk block addresses and block addresses of the virtual disks presented to the operating environment by array management software.

MB

A megabyte: an abbreviation for 1,000,000 (10 to the sixth power) bytes.

Mirroring

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

Multi-threaded

Having multiple concurrent or pseudo-concurrent execution sequences. Multi-threaded processes allow throughput-intensive applications to efficiently use resources to increase I/O performance.

Ns

A nanosecond, 10^{-9} second.

Online

An online device is a device that is accessible.

Online Expansion

Capacity expansion by adding volume or another hard drive, while the host system is accessible and/or active.

Operating Environment

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

Parity

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

Partition

A complete or partial representation of a logical drive, usually represented to a user by an operating system as a physical disk. Also called a logical volume.

PERC 4e/Di

The Dell/ITM PERC 4e/Di consists of an LSI 1030 chip on the motherboard to offer RAID control capabilities. PERC 4e/Di supports all dual-ended and LVD SCSI devices on Ultra320 and Wide SCSI channels with data transfer rates up to 320 MB/s (Megabytes per second).

PERC 4e/Di provides reliability, high performance, and fault-tolerant disk subsystem management. It is an ideal RAID solution for the internal storage of Dell's workgroup, departmental, and enterprise systems. PERC 4e/Di offers a cost-effective way to implement RAID in a server and provides reliability, high performance, and fault-tolerant disk subsystem management.

Physical Disk

A hard drive. A hard drive consists of one or more rigid magnetic discs rotating about a central axle, with associated read/write heads and electronics. A physical disk is used to store information, (data), in a non-volatile and randomly accessible memory space.

Physical Disk Roaming

The ability of adapters to detect when hard drives have been moved to different slots in the storage enclosure, such as after a hot swap.

Protocol

A set of formal rules describing how to transmit data, generally across a network or when communicating with storage sub-systems. 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.

RAID

Redundant Array of Independent Disks (originally Redundant Array of Inexpensive Disks) is an array of multiple small, independent hard drives that yields performance exceeding that of a Single Large Expensive Disk (SLED). A RAID disk subsystem can improve I/O performance relative to a system using only a single drive. The RAID array appears to the controller as a single storage unit. I/O is expedited because several disks can be accessed simultaneously. Redundant RAID levels (RAID levels 1, 5, 10, and 50), provide data protection.

RAID Level Migration

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

RAID Levels

A style of redundancy applied to an array. It can increase the performance of a logical drive though it may decrease usable capacity. Each logical array must have a RAID level assigned to it.

Read-Ahead

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

Ready State

A condition in which a workable hard drive is neither online nor a hot spare and is available to add to an array or to designate as a hot spare.

Rebuild

The regeneration of all data to a replacement disk from a failed disk in a logical drive with a RAID level 1, 5, 10 or 50 array. A disk rebuild normally occurs without interrupting normal operations on the affected logical drive, though some degradation of performance of the disk subsystem can occur.

Rebuild Rate

The percentage of CPU resources devoted to rebuilding.

Reconstruct

The act of remaking a logical drive after changing RAID levels or adding a physical drive to an existing array.

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.

Replacement Disk

A hard drive replacing a failed member disk in a RAID array.

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, adapter logic boards, power supplies and cables.

SCSI

(Small Computer System Interface) A processor-independent standard for system-level interfacing between a system and intelligent devices, such as hard disks, floppy disks, CD-ROM, printers, and scanners. SCSI can connect up to 15 devices to a single adapter (or host adapter) on the system's bus. SCSI transfers 8, 16 or 32 bits in parallel and can operate in either asynchronous or synchronous modes. The synchronous transfer rate is up to 320 MB/s.

The original standard is now called SCSI-1 to distinguish it from SCSI-2 and SCSI-3, which include specifications of Wide SCSI (a 16-bit bus) and Fast SCSI (10 MB/s transfer.) Ultra 160M SCSI is a subset of Ultra320 SCSI and allows a maximum throughput of 160 MB/s, which is twice as fast as Wide Ultra2 SCSI. Ultra320 allows a maximum throughput of 320 MB/s.

SCSI Channel

The RAID controller controls hard drives via 320M SCSI buses (channels) over which the system transfers data in either LVD or 320M SCSI modes. Each adapter controls two SCSI channels.

SCSI Disk Status

A SCSI disk drive (physical drive) can be in one of these four states:

- 1 **Online:** A powered-on and operational disk.
- 1 **Hot Spare:** A powered-on, stand-by disk ready for use if another disk fails.
- 1 **Not Responding:** The disk is not present, not powered-on, or has failed.
- 1 **Rebuild:** A disk to which one or more logical drives is restoring data.

SCSI ID

Each SCSI device on the RAID controller SCSI bus must have a different SCSI address number (Target ID or TID) from 0 to 15. Notice that one ID is used by the SCSI controller, usually ID 7. Set the SCSI ID switch on each disk drive to the correct SCSI address. See the RAID controller documentation, chassis labels or disk enclosure documentation for the correct switch settings.

Spare

A hard drive available to back up the data of other drives.

Striping

Segmentation of logically sequential data, such as a single file, so that segments can be written to multiple physical devices in a round-robin fashion. This technique is useful if the processor can read or write data faster than a single disk can supply or accept it. While data is being transferred from the first disk, the second disk can locate the next segment. Data striping is used in some modern databases and in certain RAID devices.

Stripe Size

The amount of data written to each disk. Also called "stripe depth." You can specify stripe sizes of 2 KB, 4 KB, 8 KB, 16 KB, 32 KB, 64 KB, and 128 KB for each logical drive. A larger stripe size produces improved read performance, especially if most of the reads are sequential. For mostly random reads, select a smaller stripe size.

Stripe Width

The number of disk drives across which the data are striped.

Terminator

A resistor connected to a signal wire in a bus or network for impedance matching to prevent reflections, e.g., a resistor connected across signal wires at the end of a SCSI cable.

Wide SCSI

A variant on the SCSI-2 interface. Wide SCSI uses a 16-bit bus, double the width of the original SCSI-1.

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 are written to the disk subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, elapsed time from the last cache flush, and others.

Write-Through

In Write Through caching mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction. The controller cache is not used.

[Back to Contents Page](#)

Safety Instructions

CAUTION: Safety Instructions

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

General

- 1 Do not attempt to service the computer yourself unless you are a trained service technician. Always follow installation instructions closely.
- 1 To help prevent electric shock, plug the computer and device power cables into properly grounded electrical outlets. These cables are equipped with 3-prong plugs to help ensure proper grounding. Do not use adapter plugs or remove the grounding prong from a cable. If you must use an extension cable, use a 3-wire cable with properly grounded plugs.
- 1 To help avoid the potential hazard of electric shock, do not use your computer during an electrical storm.
- 1 To help avoid the potential hazard of electric shock, do not connect or disconnect any cables or perform maintenance or reconfiguration of this product during an electrical storm.
- 1 If your computer includes a modem, the cable used with the modem should be manufactured with a minimum wire size of 26 American wire gauge (AWG) and an FCC-compliant RJ-11 modular plug.
- 1 Before you clean your computer, disconnect the computer from the electrical outlet. Clean your computer with a soft cloth dampened with water. Do not use liquid or aerosol cleaners, which may contain flammable substances.
- 1 To help avoid possible damage to the system board, wait 5 seconds after turning off the computer before disconnecting a device from the computer.
- 1 To avoid shorting out your computer when disconnecting a network cable, first unplug the cable from the network adapter on the back of your computer, and then from the network jack. When reconnecting a network cable to your computer, first plug the cable into the network jack, and then into the network adapter.
- 1 To help protect your computer from sudden, transient increases and decreases in electrical power, use a surge suppressor, line conditioner, or uninterruptible power supply (UPS).
- 1 Ensure that nothing rests on your computer's cables and that the cables are not located where they can be stepped on or tripped over.
- 1 Do not push any objects into the openings of your computer. Doing so can cause fire or electric shock by shorting out interior components.
- 1 Keep your computer away from radiators and heat sources. Also, do not block cooling vents. Avoid placing loose papers underneath your computer; do not place your computer in a closed-in wall unit or on a bed, sofa, or rug.

When Using Your Computer

As you use your computer, observe the following safe-handling guidelines.

CAUTION: Do not operate your computer with any cover(s) (including computer covers, bezels, filler brackets, front-panel inserts, and so on) removed.

Your computer is equipped with one of the following:

- 1 A fixed-voltage power supply — Computers with a fixed-voltage power supply do not have a voltage selection switch on the back panel and operate at only one voltage (see the regulatory label on the outside of the computer for its operating voltage).
- 1 An auto-sensing voltage circuit — Computers with an auto-sensing voltage circuit do not have a voltage selection switch on the back panel and automatically detect the correct operating voltage.
- 1 A manual voltage selection switch — Computers with a voltage selection switch on the back panel must be manually set to operate at the correct operating voltage.

To help avoid damaging a computer with a manual voltage selection switch, ensure that the voltage selection switch is set to match the AC power available at your location:

- 1 115 V/60 Hz in most of North and South America and some Far Eastern countries such as South Korea and Taiwan

 **NOTE:** The voltage selection switch must be set to the 115-V position even though the AC power available in Japan is 100 V.

- 1 100 V/50 Hz in eastern Japan and 100 V/60 Hz in western Japan
- 1 230 V/50 Hz in some regions in the Caribbean and South America and most of Europe, the Middle East, and the Far East

Also, ensure that your monitor and attached devices are electrically rated to operate with the AC power available in your location.

- 1 Before working inside the computer, unplug the computer to help prevent electric shock or system board damage. Certain system board components continue to receive power any time the computer is connected to AC power.

When Working Inside Your Computer

Before you open the computer cover, perform the following steps in the sequence indicated.

 **CAUTION:** Do not attempt to service the computer yourself, except as explained in your online Dell™ documentation or in instructions otherwise provided to you by Dell. Always follow installation and service instructions closely.

 **NOTICE:** To help avoid possible damage to the system board, wait 5 seconds after turning off the computer before removing a component from the system board or disconnecting a device from the computer.

1. Perform an orderly computer shutdown using the operating system menu.
2. Turn off your computer and any devices connected to the computer.
3. Ground yourself by touching an unpainted metal surface on the chassis, such as the metal around the card-slot openings at the back of the computer, before touching anything inside your computer.

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

4. Disconnect your computer and devices, including the monitor, from their electrical outlets. Also, disconnect any telephone or telecommunication lines from the computer.

Doing so reduces the potential for personal injury or shock.

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

- 1 When you disconnect a cable, pull on its connector or on its strain-relief loop, not on the cable itself. Some cables have a connector with locking tabs; if you are disconnecting this type of cable, press in on the locking tabs before disconnecting the cable. As you pull connectors apart, keep them evenly aligned to avoid bending any connector pins. Also, before you connect a cable, ensure that both connectors are correctly oriented and aligned.
- 1 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.

 **CAUTION:** There is a danger of a new battery exploding if it is incorrectly installed. Replace the battery only with the same or equivalent type recommended by the manufacturer. 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.

Protecting Against Electrostatic Discharge

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

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

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

- 1 Do not remove components from their antistatic packing material until you are ready to install the component in your computer. Just before unwrapping the antistatic packaging, discharge static electricity from your body.
- 1 When transporting an electrostatic sensitive component, first place it in an antistatic container or packaging.
- 1 Handle all electrostatic sensitive components in a static-safe area. If possible, use antistatic floor pads and workbench pads.

Ergonomic Computing Habits

 **CAUTION:** Improper or prolonged keyboard use may result in injury.

 **CAUTION:** Viewing the monitor screen for extended periods of time may result in eye strain.

Battery Disposal

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.

[Back to Contents Page](#)