

User's Guide



# LSI Logic MegaRAID®

SAS Software

English



# **USER'S GUIDE**

## **MegaRAID SAS Software**

**April 2009**

---

This document contains proprietary information of LSI Corporation. The information contained herein is not to be used by or disclosed to third parties without the express written permission of an officer of LSI Corporation.

Document 80-00156-01 Rev. E (December 2008)

This document describes the LSI Corporation's MegaRAID Storage Manager software. This document will remain the official reference source for all revisions/releases of this product until rescinded by an update.

LSI Corporation reserves the right to make changes to any products herein at any time without notice. LSI does not assume any responsibility or liability arising out of the application or use of any product described herein, except as expressly agreed to in writing by LSI; nor does the purchase or use of a product from LSI convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI or third parties.

Copyright © 2005-2008 by LSI Corporation. All rights reserved.

#### TRADEMARK ACKNOWLEDGMENT

LSI, the LSI logo design, iBBU, MegaRAID, and MegaRAID Storage Manager are trademarks or registered trademarks of LSI Corporation. Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Intel and Pentium are registered trademarks of Intel Corporation. SCO and SCO UnixWare are registered trademarks and OpenServer is a trademark of the SCO Group, Inc. This product includes software developed by the Apache Software Foundation (<http://www.apache.org>). All other brand and product names may be trademarks of their respective companies.

CD

**To receive product literature, visit us at <http://www.lsi.com>.**

**For a current list of our distributors, sales offices, and design resource centers, view our web page located at**

**<http://www.lsi.com/cm/ContactSearch.do?locale=EN>**

# Preface

---

This document explains how to use the MegaRAID Storage Manager™ software, WebBIOS, and Command Line Interface (CLI) utilities to configure, monitor, and maintain MegaRAID® Serial-attached SCSI (SAS) RAID controllers and the storage-related devices connected to them.

---

## Audience

This document assumes that you are familiar with SAS controllers and configuration utilities. The people who benefit from this book are network administrators who need to create storage configurations on LSI SAS controllers.

---

## Organization

This document has the following chapters and appendixes:

- [Chapter 1, “Overview,”](#) describes the SAS, Serial ATA (SATA) II, and Solid State Disk (SSD) technologies, configuration scenarios, and Technical Support information.
- [Chapter 2, “Introduction to RAID,”](#) describes RAID (Redundant Array of Independent Disks), RAID functions and benefits, RAID components, RAID levels, and configuration strategies. In addition, it defines the RAID availability concept, and offers tips for configuration planning.
- [Chapter 3, “WebBIOS Configuration Utility,”](#) explains how to use the pre-boot WebBIOS Configuration Utility to create and manage storage configurations.

- [Chapter 4, “MegaRAID Command Tool,”](#) explains how to use the MegaRAID Command Tool to create and manage storage configurations. The MegaRAID Command Tool is a CLI application for SAS.
- [Appendix, “Glossary,”](#) contains definitions of storage-related terms.

## Conventions

**Note:** Notes contain supplementary information that can have an effect on system performance.

**Caution:** Cautions are notifications that an action has the potential to adversely affect equipment operation, system performance, or data integrity.

## Revision History

Document Number	Date Revision	Remarks
80-00156-01 Rev. E	December 2008	Added the overview chapter. Updated the WebBIOS Configuration Utility, MegaRAID Storage Manager, and MegaCLI chapters.
80-00156-01 Rev. D	April 2008	Updated the RAID overview section. Updated the WebBIOS Configuration Utility and the MegaRAID Storage Manager. Updated the MegaCLI commands.
80-00156-01 Rev. C	July 2007 Version 2.1	Updated operating system support for MegaCLI.
80-00156-01 Rev. B	June 2007 Version 2.0	Updated the WebBIOS Configuration Utility and the MegaRAID Storage Manager. Updated the MegaCLI commands. Added the RAID introduction chapter.
80-00156-01 Rev. A	August 2006 Version 1.1	Corrected the procedure for creating RAID 10 and RAID 50 drive groups in the WebBIOS Configuration Utility.
DB15-000339-00	December 2005 Version 1.0	Initial release of this document.

# Contents

---

## Chapter 1 Overview

1.1	SAS Technology	1-1
1.2	Serial-attached SCSI Device Interface	1-3
1.3	Serial ATA II Features	1-3
1.4	Solid State Disk Features	1-4
1.5	Configuration Scenarios	1-5
1.5.1	Valid Drive Mix Configurations with HDDs and SSDs	1-7

---

## Chapter 2 Introduction to RAID

2.1	RAID Description	2-1
2.2	RAID Benefits	2-1
2.3	RAID Functions	2-1
2.4	Components and Features	2-2
2.4.1	Physical Array	2-2
2.4.2	Virtual Drive	2-2
2.4.3	RAID Drive Group	2-3
2.4.4	Fault Tolerance	2-3
2.4.5	Consistency Check	2-5
2.4.6	Copyback	2-5
2.4.7	Background Initialization	2-6
2.4.8	Patrol Read	2-7
2.4.9	Disk Striping	2-7
2.4.10	Disk Mirroring	2-8
2.4.11	Parity	2-9
2.4.12	Disk Spanning	2-10
2.4.13	Hot Spares	2-11

2.4.14	Disk Rebuilds	2-13
2.4.15	Rebuild Rate	2-14
2.4.16	Hot Swap	2-14
2.4.17	Drive States	2-14
2.4.18	Virtual Drive States	2-15
2.4.19	Enclosure Management	2-16
2.5	RAID Levels	2-16
2.5.1	Summary of RAID Levels	2-16
2.5.2	Selecting a RAID Level	2-17
2.5.3	RAID 0	2-18
2.5.4	RAID 00	2-19
2.5.5	RAID 1	2-20
2.5.6	RAID 5	2-21
2.5.7	RAID 6	2-22
2.5.8	RAID 10	2-23
2.5.9	RAID 50	2-25
2.5.10	RAID 60	2-26
2.6	RAID Configuration Strategies	2-28
2.6.1	Maximizing Fault Tolerance	2-28
2.6.2	Maximizing Performance	2-30
2.6.3	Maximizing Storage Capacity	2-32
2.7	RAID Availability	2-33
2.7.1	RAID Availability Concept	2-33
2.8	Configuration Planning	2-35
2.8.1	Number of Drives	2-35
2.8.2	Drive Group Purpose	2-35

---

## Chapter 3

### WebBIOS Configuration Utility

3.1	Overview	3-1
3.2	Starting the WebBIOS CU	3-2
3.3	WebBIOS CU Main Screen Options	3-3
3.4	Creating a Storage Configuration	3-6
3.4.1	Selecting the Configuration with the Configuration Wizard	3-6
3.4.2	Using Automatic Configuration	3-9
3.4.3	Using Manual Configuration	3-9

3.5	Viewing and Changing Device Properties	3-60
3.5.1	Viewing and Changing Controller Properties	3-60
3.5.2	Viewing and Changing Virtual Drive Properties	3-64
3.5.3	Viewing Drive Properties	3-66
3.5.4	Viewing Battery Backup Unit Information	3-68
3.6	Viewing System Event Information	3-70
3.7	Managing Configurations	3-71
3.7.1	Running a Consistency Check	3-71
3.7.2	Deleting a Virtual Drive	3-72
3.7.3	Importing or Clearing a Foreign Configuration	3-72
3.7.4	Migrating the RAID Level of a Virtual Drive	3-77

---

## Chapter 4

### MegaRAID Command Tool

4.1	Product Overview	4-2
4.2	Novell NetWare, SCO, Solaris, FreeBSD, and DOS Operating System Support	4-4
4.3	Command Line Abbreviations and Conventions	4-4
4.3.1	Abbreviations Used in the Command Line	4-4
4.3.2	Conventions	4-5
4.4	Controller Property-Related Options	4-6
4.4.1	Display Controller Properties	4-7
4.4.2	Display Number of Controllers Supported	4-7
4.4.3	Enable or Disable Automatic Rebuild	4-7
4.4.4	Flush Controller Cache	4-8
4.4.5	Set Controller Properties	4-8
4.4.6	Display Specified Controller Properties	4-9
4.4.7	Set Factory Defaults	4-10
4.4.8	Set SAS Address	4-10
4.4.9	Set Time and Date on Controller	4-10
4.4.10	Display Time and Date on Controller	4-10
4.5	Patrol Read-Related Controller Properties	4-11
4.5.1	Set Patrol Read Options	4-11
4.5.2	Set Patrol Read Delay Interval	4-12
4.6	BIOS-Related Properties	4-12
4.6.1	Set or Display Bootable Virtual Drive ID	4-12
4.6.2	Select BIOS Status Options	4-13

4.7	Battery Backup Unit-Related Properties	4-13
4.7.1	Display BBU Information	4-13
4.7.2	Display BBU Status Information	4-14
4.7.3	Display BBU Capacity	4-15
4.7.4	Display BBU Design Parameters	4-15
4.7.5	Display Current BBU Properties	4-16
4.7.6	Start BBU Learning Cycle	4-16
4.7.7	Place Battery in Low-Power Storage Mode	4-16
4.7.8	Set BBU Properties	4-17
4.8	Options for Displaying Logs Kept at Firmware Level	4-17
4.8.1	Event Log Management	4-17
4.8.2	Set BBU Terminal Logging	4-18
4.9	Configuration-Related Options	4-19
4.9.1	Create a RAID Drive Group from All Unconfigured Good Drives	4-19
4.9.2	Add RAID 0, 1, 5, or 6 Configuration	4-20
4.9.3	Add RAID 10, 50, or 60 Configuration	4-22
4.9.4	Clear the Existing Configuration	4-22
4.9.5	Save the Configuration on the Controller	4-23
4.9.6	Restore the Configuration Data from File	4-23
4.9.7	Manage Foreign Configuration Information	4-23
4.9.8	Delete Specified Virtual Drive(s)	4-24
4.9.9	Display the Free Space	4-24
4.10	Virtual Drive-Related Options	4-25
4.10.1	Display Virtual Drive Information	4-25
4.10.2	Change the Virtual Drive Cache and Access Parameters	4-25
4.10.3	Display the Virtual Drive Cache and Access Parameters	4-26
4.10.4	Manage Virtual Drives Initialization	4-26
4.10.5	Manage a Consistency Check	4-27
4.10.6	Manage a Background Initialization	4-27
4.10.7	Perform a Virtual Drive Reconstruction	4-28
4.10.8	Display Information about Virtual Drives and Drives	4-28
4.10.9	Display the Number of Virtual Drives	4-29

4.11	Drive-Related Options	4-29
4.11.1	Display Drive Information	4-29
4.11.2	Set the Drive State to Online	4-30
4.11.3	Set the Drive State to Offline	4-30
4.11.4	Change the Drive State to Unconfigured Good	4-30
4.11.5	Change Drive State	4-31
4.11.6	Manage a Drive Initialization	4-31
4.11.7	Rebuild a Drive	4-32
4.11.8	Locate the Drive(s) and Activate LED	4-32
4.11.9	Mark the Configured Drive as Missing	4-32
4.11.10	Display the Drives in Missing Status	4-33
4.11.11	Replace the Configured Drives and Start an Automatic Rebuild	4-33
4.11.12	Prepare the Unconfigured Drive for Removal	4-33
4.11.13	Display Total Number of Drives	4-34
4.11.14	Display List of Physical Devices	4-34
4.11.15	Download Firmware to the Physical Devices	4-34
4.12	Enclosure-Related Options	4-35
4.13	Flashing the Firmware	4-35
4.13.1	Flash the Firmware with the ROM File	4-35
4.13.2	Flash the Firmware in Mode 0 with the ROM File	4-36
4.14	SAS Topology	4-36
4.15	Diagnostic-Related Options	4-37
4.15.1	Start Controller Diagnostics	4-37
4.15.2	Start Battery Test	4-37
4.15.3	Start NVRAM Diagnostic	4-37
4.16	Miscellaneous Options	4-38
4.16.1	Display the MegaCLI Version	4-38
4.16.2	Display Help for MegaCLI	4-38

---

**Appendix**  
**Glossary**



---

## Figures

1.1	Example of an LSI SAS Direct-Connect Application	1-6
1.2	Example of an LSI SAS RAID Controller Configured with an LSISASx12 Expander	1-6
2.1	Example of Disk Striping (RAID 0)	2-7
2.2	Example of Disk Mirroring (RAID 1)	2-8
2.3	Example of Distributed Parity (RAID 5)	2-9
2.4	Example of Disk Spanning	2-10
2.5	RAID 0 Drive Group Example with Two Drives	2-18
2.6	RAID 00 Drive Group Example with Two Drives	2-20
2.7	RAID 1 Drive Group	2-21
2.8	RAID 5 Drive Group with Six Drives	2-22
2.9	Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)	2-23
2.10	RAID 10 Level Virtual Drive	2-25
2.11	RAID 50 Level Virtual Drive	2-26
2.12	RAID 60 Level Virtual Drive	2-28
3.1	WebBIOS CU Main Screen	3-3
3.2	WebBIOS Configuration Wizard Screen	3-6
3.3	WebBIOS Configuration Method Screen	3-7
3.4	WebBIOS Drive Group Definition Screen	3-10
3.5	WebBIOS Virtual Drive Definition Screen	3-11
3.6	RAID 0 Configuration Preview	3-14
3.7	WebBIOS Drive Group Definition Screen	3-16
3.8	WebBIOS Span Definition Screen	3-17
3.9	WebBIOS Virtual Drive Definition Screen	3-18
3.10	RAID 00 Configuration Preview	3-21
3.11	WebBIOS Disk Group Definition Screen	3-22
3.12	WebBIOS Virtual Drive Definition Screen	3-23
3.13	RAID 1 Configuration Preview	3-26
3.14	WebBIOS Disk Group Definition Screen	3-27
3.15	WebBIOS Virtual Drive Definition Screen	3-28
3.16	RAID 5 Configuration Preview	3-31
3.17	WebBIOS Drive Group Definition Screen	3-33
3.18	WebBIOS Virtual Drive Definition Screen	3-34
3.19	RAID 6 Configuration Preview	3-37
3.20	WebBIOS Drive Group Definition Screen	3-39

3.21	WebBIOS Span Definition Screen	3-40
3.22	WebBIOS Virtual Drive Definition Screen	3-41
3.23	RAID 10 Configuration Preview	3-44
3.24	WebBIOS Disk Group Definition Screen	3-46
3.25	WebBIOS Span Definition Screen	3-47
3.26	WebBIOS Virtual Drive Definition Screen	3-48
3.27	RAID 50 Configuration Preview	3-51
3.28	WebBIOS Drive Group Definition Screen	3-53
3.29	WebBIOS Span Definition Screen	3-54
3.30	WebBIOS Virtual Drive Definition Screen	3-55
3.31	RAID 60 Configuration Preview	3-58
3.32	First Controller Properties Screen	3-59
3.33	Second Controller Properties Screen	3-60
3.34	Third Controller Properties Screen	3-61
3.35	Virtual Drive Screen	3-64
3.36	Physical Drive Screen	3-66
3.37	Battery Module Screen	3-68
3.38	Event Information Screen	3-69
3.39	Foreign Configuration Import Screen	3-72
3.40	Foreign Configuration Preview Screen	3-73

---

## Tables

1.1	Valid Drive Mix Configurations	1-7
2.1	Types of Parity	2-9
2.2	Spanning for RAID 10, RAID 50, and RAID 60	2-11
2.3	Drive States	2-14
2.4	Virtual Drive States	2-15
2.5	RAID 0 Overview	2-18
2.6	RAID 00 Overview	2-19
2.7	RAID 1 Overview	2-20
2.8	RAID 5 Overview	2-21
2.9	RAID 6 Overview	2-23
2.10	RAID 10 Overview	2-24
2.11	RAID 50 Overview	2-25
2.12	RAID 60 Overview	2-27
2.13	RAID Levels and Fault Tolerance	2-29
2.14	RAID Levels and Performance	2-30
2.15	RAID Levels and Capacity	2-32
2.16	Factors to Consider for Drive Group Configuration	2-36
3.1	WebBIOS CU Toolbar Icons	3-4
3.2	Controller Properties Menu Options	3-61
3.3	Additional Drives Required for RAID-Level Migration	3-76
4.1	Command Line Abbreviations	4-5
4.2	Conventions	4-5
4.3	Controller Parameters	4-6
4.4	Number of Controllers Supported	4-7
4.5	Enable or Disable Automatic Rebuild	4-7
4.6	Cache Flush on Selected Controller	4-7
4.7	Set Controller Properties	4-8
4.8	Display Specified Controller Properties	4-9
4.9	Set Factory Defaults	4-10
4.10	Set SAS Address on Controller	4-10
4.11	Set Time and Date on Controller	4-10
4.12	Display Time and Date on Controller	4-10
4.13	Set Patrol Read Options	4-11
4.14	Set Patrol Read Delay Interval	4-12
4.15	Bootable Virtual Drive ID	4-12
4.16	Options for BIOS Status	4-12

4.17	Display BBU Information	4-13
4.18	Display BBU Status Information	4-13
4.19	Display BBU Capacity Information	4-15
4.20	Display BBU Design Parameters	4-15
4.21	Display Current BBU Properties	4-16
4.22	Start BBU Learning Cycle	4-16
4.23	Place Battery in Low-Power Storage Mode	4-16
4.24	Set BBU Properties	4-17
4.25	Event Log Management	4-17
4.26	Set BBU Terminal Logging	4-18
4.27	Create a Drive Group from All of the Unconfigured Drives	4-19
4.28	Add RAID 0, 1, 5, or 6 Configuration	4-20
4.29	Add RAID 10, 50, or 60 Configuration	4-21
4.30	Clear Existing Configuration	4-21
4.31	Save Configuration on the Controller	4-22
4.32	Restore Configuration Data from File	4-22
4.33	Manage Foreign Configuration Information	4-23
4.34	Delete Specified Virtual Drives	4-23
4.35	Display Free Space	4-23
4.36	Display Virtual Drive Information	4-24
4.37	Change Virtual Drive Cache and Access Parameters	4-24
4.38	Display Virtual Drive Cache and Access Parameters	4-25
4.39	Manage Virtual Drive Initialization	4-25
4.40	Manage Consistency Check	4-26
4.41	Manage Background Initialization	4-26
4.42	Virtual Drive Reconstruction	4-27
4.43	Display Virtual Drive and Drive Information	4-27
4.44	Display Number of Virtual Drives	4-27
4.45	Display Drive Information	4-28
4.46	Set Drive State to Online	4-28
4.47	Set Drive State to Offline	4-29
4.48	Change Drive State to Unconfigured Good	4-29
4.49	Change Drive State	4-29
4.50	Drive Initialization	4-30
4.51	Rebuild a Drive	4-30
4.52	Locate Drive and Activate LED	4-31
4.53	Mark Configured Drive as Missing	4-31
4.54	Display Drives in Missing Status	4-31

4.55	Replace Configured Drive(s) and Start Automatic Rebuild	4-31
4.56	Prepare Unconfigured Drive(s) for Removal	4-32
4.57	Display Number of Drives Attached to an Controller	4-32
4.58	Display List of Physical Devices Attached to Controller(s)	4-32
4.59	Download Firmware to the Physical Devices	4-33
4.60	Display Enclosure Information	4-33
4.61	Flash Firmware with ROM File	4-33
4.62	Flash Firmware in Mode 0 with ROM File	4-34
4.63	Display PHY Connection Information	4-35
4.64	Start Diagnostics Setting	4-35
4.65	Start Battery Test	4-35
4.66	Start NVRAM Diagnostic	4-36
4.67	Display MegaCLI Version	4-36
4.68	Display Help for MegaCLI	4-36



# Chapter 1

## Overview

---

This guide documents the utilities used to configure, monitor, and maintain MegaRAID® Serial-attached SCSI (SAS) RAID controllers with RAID control capabilities and the storage-related devices connected to them. This guide explains how to use the MegaRAID Storage Manager™ software, WebBIOS, and Command Line Interface (CLI). In addition, it documents SAS technology, Serial ATA (SATA) technology, Solid State Disk (SSD) technology, configuration scenarios, and drive types.

This chapter consists of the following sections:

- [Section 1.1, “SAS Technology”](#)
  - [Section 1.2, “Serial-attached SCSI Device Interface”](#)
  - [Section 1.3, “Serial ATA II Features”](#)
  - [Section 1.4, “Solid State Disk Features”](#)
  - [Section 1.5, “Configuration Scenarios”](#)
- 

## 1.1 SAS Technology

The MegaRAID SAS RAID controllers are high-performance intelligent PCI Express-to-SCSI/Serial ATA II controllers with RAID control capabilities. MegaRAID SAS RAID controllers provide reliability, high performance, and fault-tolerant disk subsystem management. They are an ideal RAID solution for the internal storage of workgroup, departmental, and enterprise systems. MegaRAID SAS RAID controllers offer a cost-effective way to implement RAID in a server.

SAS technology brings a wealth of options and flexibility with the use of SAS devices, Serial ATA (SATA) II devices, and SSD devices within the same storage infrastructure. These devices bring individual characteristics that make each one a more suitable choice depending on

your storage needs. MegaRAID gives you the flexibility to combine these two similar technologies on the same controller, within the same enclosure, and in the same virtual drive.

**Note:** LSI recommends that you carefully assess any decision to mix SAS drives and SATA drives within the same *virtual drives*. Although you can mix drives, LSI strongly discourages the practice. This recommendation applies to both HDDs and SSDs.

The MegaRAID SAS RAID controllers are based on the LSI first-to-market SAS IC technology and proven MegaRAID technology. As second-generation PCI Express RAID controllers, the MegaRAID SAS RAID controllers address the growing demand for increased data throughput and scalability requirements across midrange and enterprise-class server platforms. LSI offers a family of MegaRAID SAS RAID controllers addressing the needs for both internal and external solutions.

The SAS controllers support the ANSI *Serial Attached SCSI standard, version 1.1*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, version 1.0a*. Supporting both the SAS and SATA II interfaces, the SAS controller is a versatile controller that provides the backbone of both server environments and high-end workstation environments.

Each port on the SAS RAID controller supports SAS devices, SATA II devices, or SSD devices using the following protocols:

- SAS Serial SCSI Protocol (SSP), which enables communication with other SAS devices
- SATA II, which enables communication with other SATA II devices
- Serial Management Protocol (SMP), which communicates topology management information directly with an attached SAS expander device
- Serial Tunneling Protocol (STP), which enables communication with a SATA II device through an attached expander

---

## 1.2 Serial-attached SCSI Device Interface

SAS is a serial, point-to-point, enterprise-level device interface that leverages the proven SCSI protocol set. SAS is a convergence of the advantages of SATA II, SCSI, and Fibre Channel, and is the future mainstay of the enterprise and high-end workstation storage markets. SAS offers a higher bandwidth per pin than parallel SCSI, and it improves signal and data integrity.

The SAS interface uses the proven SCSI command set to ensure reliable data transfers, while providing the connectivity and flexibility of point-to-point serial data transfers. The serial transmission of SCSI commands eliminates clock-skew challenges. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.

SAS controllers leverage a common electrical and physical connection interface that is compatible with Serial ATA technology. The SAS and SATA II protocols use a thin, 7-wire connector instead of the 68-wire SCSI cable or 26-wire ATA cable. The SAS/SATA II connector and cable are easier to manipulate, allow connections to smaller devices, and do not inhibit airflow. The point-to-point SATA II architecture eliminates inherent difficulties created by the legacy ATA master-slave architecture, while maintaining compatibility with existing ATA firmware.

---

## 1.3 Serial ATA II Features

The SATA bus is a high-speed, internal bus that provides a low pin count, low voltage level bus for device connections between a host controller and a SATA device.

The following list describes the SATA II features of the RAID controllers:

- Supports SATA II data transfers of 3.0 Gbits/s
- Supports STP data transfers of 3.0 Gbits/s
- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices

- Eliminates the master-slave construction used in parallel ATA
- Allows addressing of multiple SATA II targets through an expander
- Allows multiple initiators to address a single target (in a fail-over configuration) through an expander

---

## 1.4 Solid State Disk Features

MegaRAID firmware supports SSD drives attached to MegaRAID SAS controllers. These drives are expected to behave like SATA or SAS HDD drives. The major advantages of SSD drives include:

- High random read speed (because there is no read-write head to move)
- High performance-to-power ratio, as these drives have very low power consumption compared to HDDs
- Low latency
- High mechanical reliability
- Lower weight and size (for low-capacity SSD drives)

The features and operations on SSD drives are the same as for hard disk drives (HDD).

**Note:** MegaRAID implements support for only those SATA SSD drives which support ATA-8 ACS compliance.

You can choose whether to allow a virtual drive to consist of both SSD devices and HDDs. For a virtual drive that consists of SSDs only, you can choose whether to allow SAS SSD drives and SATA SSD drives in that virtual drive. For virtual drives that have both SSDs and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA SSD devices in various combinations.

**Note:** Support for SATA SDD drives applies only to those drives that support ATA-8 ACS compliance.

---

## 1.5 Configuration Scenarios

There are three main scenarios in which you can use the SAS RAID controllers:

- **Low-end, internal SATA II configurations:** In this configuration, use the RAID controller as a high-end SATA II compatible controller that connects up to eight disks either directly or through a port expander. This configuration is mostly for low-end or entry servers. Enclosure management is provided through out-of-band I<sup>2</sup>C bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.
- **Midrange internal SAS configurations:** This configuration is like the internal SATA II configurations, but with high-end disks. This configuration is more suitable for low-range to midrange servers.
- **High-end external SAS/SATA II configurations:** This configuration is for both internal connectivity and external connectivity, using SATA II drives, SAS drives, or both. External enclosure management is supported through in-band, SCSI-enclosed storage. The configuration must support STP and SMP.

Figure 1.1 shows a direct-connect configuration. The Inter-IC (I<sup>2</sup>C) interface communicates with peripherals. The external memory bus provides a 32-bit memory bus, parity checking, and chip select signals for pipelined synchronous burst static random access memory (PSBRAM), nonvolatile static random access memory (NVSRAM), and Flash ROM.

Note: The external memory bus is 32-bit for the SAS 8704ELP and the SAS 8708ELP, and 64-bit for the SAS 8708EM2, the SAS 8880EM2, and the SAS 8888ELP.

**Figure 1.1 Example of an LSI SAS Direct-Connect Application**

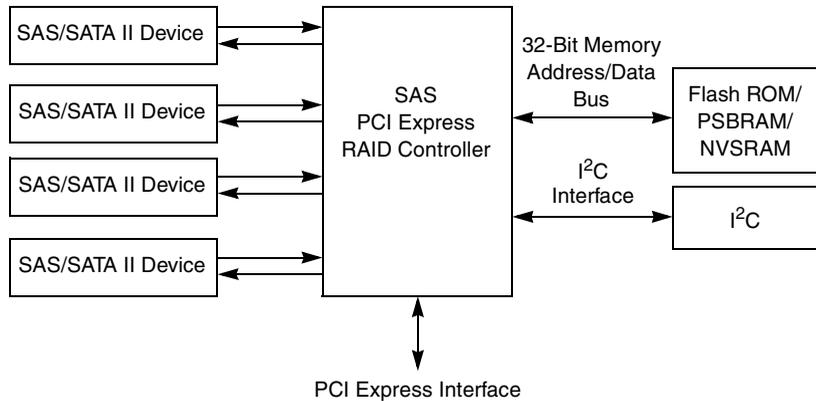
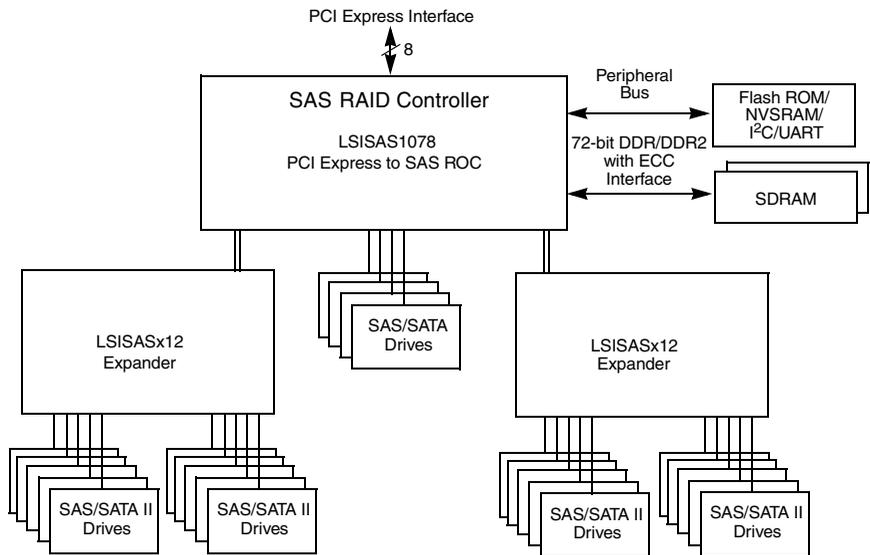


Figure 1.2 shows an example of a SAS RAID controller configured with an LSI SASx12 expander that is connected to SAS disks, SATA II disks, or both.

**Figure 1.2 Example of an LSI SAS RAID Controller Configured with an LSI SASx12 Expander**



## 1.5.1 Valid Drive Mix Configurations with HDDs and SSDs

You can allow a virtual drive to consist of both SSDs and HDDs. For virtual drives that have both SSDs and HDDs, you can choose whether to mix SAS drives and SATA drives on the SSD devices.

You can choose whether to allow a virtual drive to consist of both SSD devices and HDDs. For a virtual drive that consists of SSDs only, you can choose whether to allow SAS SSD drives and SATA SSD drives in that virtual drive. For virtual drives that have both SSDs and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA SSD devices in various combinations.

[Table 1.1](#) lists the valid drive mix configurations you can use when you create virtual drives and allow HDD and SSD mixing. The valid drive mix configurations are based on manufacturer settings.

**Table 1.1 Valid Drive Mix Configurations**

#	Valid Drive Mix Configurations
1.	SAS HDD with SAS SDD (SAS-only configuration)
2.	SATA HDD with SATA SSD (SATA-only configuration)
3.	SAS HDD with a mix of SAS and SATA SSD (a SATA HDD cannot be added)
4.	SATA HDD with a mix of SAS and SATA SSD (a SAS HDD cannot be added)
5.	SAS SSD with a mix of SAS and SATA HDD (a SATA SSD cannot be added)
6.	SATA SSD with a mix of SAS and SATA HDD (a SAS SSD cannot be added)
7.	A mix of SAS and SATA HDD with a mix of SAS and SATA SSD
8.	A SSD cannot be added to a HDD, but a SAS/SATA mix is allowed.

**Note:** Only one of the valid configurations listed in [Table 1.1](#) is allowed based on your controller card manufacturing setting.

**Note:** The valid drive mix also applies to hot spares. For hot spare information, see [Section 2.4.13, “Hot Spares,” page 2-11](#).



# Chapter 2

## Introduction to RAID

---

This chapter describes RAID (Redundant Array of Independent Disks), RAID functions and benefits, RAID components, RAID levels, and configuration strategies. In addition, it defines the RAID availability concept, and offers tips for configuration planning.

---

### 2.1 RAID Description

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

---

### 2.2 RAID Benefits

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

---

### 2.3 RAID Functions

Virtual drives are drive groups or spanned drive groups that are available to the operating system. The storage space in a virtual drive is spread across all of the drives in the drive group.

Your drives must be organized into virtual drives in a drive group and they must be able to support the RAID level that you select. Below are some common RAID functions:

- Creating hot spare drives
- Configuring drive groups and virtual drives
- Initializing one or more virtual drives
- Accessing controllers, virtual drives, and drives individually
- Rebuilding failed drives
- Verifying that the redundancy data in virtual drives using RAID level 1, 5, 6, 10, 50, or 60 is correct
- Reconstructing virtual drives after changing RAID levels or adding a drive to a drive group
- Selecting a host controller to work on

---

## 2.4 Components and Features

RAID levels describe a system for ensuring the availability and redundancy of data stored on large disk subsystems. See [Section 2.5, “RAID Levels,”](#) for detailed information about RAID levels. The following subsections describes the components of RAID drive groups and RAID levels.

### 2.4.1 Physical Array

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

### 2.4.2 Virtual Drive

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

### 2.4.3 RAID Drive Group

A RAID drive group is one or more drives controlled by the RAID controller.

### 2.4.4 Fault Tolerance

Fault tolerance is the capability of the subsystem to undergo a drive failure or failures without compromising data integrity, and processing capability. The RAID controller provides this support through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. The system can still work properly even with drive failure in a drive group, though performance can be degraded to some extent.

In a span of RAID 1 drive groups, each RAID 1 drive group has two drives and can tolerate one drive failure. The span of RAID 1 drive groups can contain up to 32 drives, and tolerate up to 16 drive failures - one in each drive group. A RAID 5 drive group can tolerate one drive failure in each RAID 5 drive group. A RAID 6 drive group can tolerate up to two drive failures.

Each spanned RAID 10 virtual drive can tolerate multiple drive failures, as long as each failure is in a separate drive group. A RAID 50 virtual drive can tolerate two drive failures, as long as each failure is in a separate drive group. RAID 60 drive groups can tolerate up to two drive failures in each drive group.

**Note:** RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual 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 that it is also important for the system to be available during the repair of the problem.

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

#### **2.4.4.1 Multipathing**

The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.

Applications show the enclosures and the drives connected to the enclosures. The firmware dynamically recognizes new enclosures added to a configuration along with their contents (new drives). In addition, the firmware dynamically adds the enclosure and its contents to the management entity currently in-use.

Multipathing provides the following features:

- Support for failover, in the event of path failure
- Auto-discovery of new or restored paths while the system is online, and reversion to system load balancing policy
- Measurable bandwidth improvement to the multi-path device
- Support for changing the load balancing path while the system is online

The firmware determines whether enclosure modules (ESMs) are part of the same enclosure. When a new enclosure module is added (allowing multi-path) or removed (going single path), an Asynchronous Event Notification (AEN) is generated. AENs about drives contain correct information about the "enclosure", when the drives are connected by multiple paths. The enclosure module detects partner ESMs and issue events appropriately.

In a system with two ESMs, you can replace one of the ESMs without affecting the virtual drive availability. For example, the controller can run heavy I/Os, and when you replace one of the ESM modules, I/Os should

not stop. The controller uses different paths to balance the load on the entire system.

In the MegaRAID Storage Manager utility, when multiple paths are available to a drive, the drive information will show only one enclosure. The utility shows that a redundant path is available to a drive. All drives with a redundant path display this information. The firmware supports online replacement of enclosure modules.

## 2.4.5 Consistency Check

The Consistency Check operation verifies correctness of the data in virtual drives that use RAID levels 1, 5, 6, 10, 50, and 60. (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.

## 2.4.6 Copyback

The copyback feature allows you to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. Copyback is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). Copyback can be run automatically or manually.

Typically, when a drive fails or is expected to fail, the data is rebuilt on a hot spare. The failed drive is replaced with a new disk. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.

Copyback is also initiated when the first Self-Monitoring Analysis and Reporting Technology (SMART) error occurs on a drive that is part of a virtual drive. The destination drive is a hot spare that qualifies as a rebuild drive. The drive with the SMART error is marked as "failed" only after the successful completion of the copyback. This avoids putting the drive group in degraded status.

**Note:** During a copyback operation, if the drive group involved in the copyback is deleted because of a virtual drive deletion, the destination drive reverts to an Unconfigured Good state or hot spare state.

### **Order of Precedence –**

In the following scenarios, rebuild takes precedence over the copyback operation:

1. If a copyback operation is already taking place to a hot spare drive, and any virtual drive on the controller degrades, the copyback operation aborts, and a rebuild starts. The rebuild changes the virtual drive to the optimal state.
2. The rebuild operation takes precedence over the copyback operation when the conditions exist to start both operations. For example:
  - a. Where the hot spare is not configured (or unavailable) in the system.
  - b. There are two drives (both members of virtual drives), with one drive exceeding the SMART error threshold, and the other failed.
  - c. If you add a hot spare (assume a global hot spare) during a copyback operation, the copyback is aborted, and the rebuild operation starts on the hot spare.

## **2.4.7 Background Initialization**

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

Background initialization is a check for media errors on the drives. It ensures that striped data segments are the same on all drives in a drive group. The default and recommended background initialization rate is 30 percent. 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.

## 2.4.8 Patrol Read

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

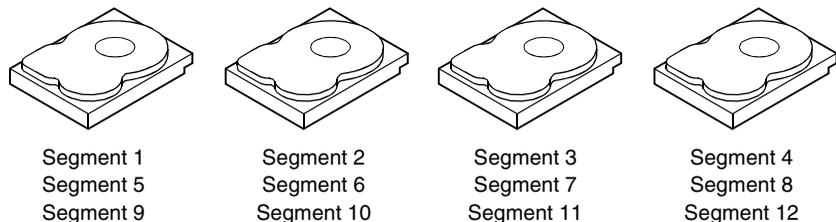
You can use the MegaRAID Command Tool or the MegaRAID Storage Manager to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See [Section 4.4, “Controller Property-Related Options,”](#).

## 2.4.9 Disk Striping

Disk striping allows you to write data across multiple drives instead of just one drive. Disk striping involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. It is recommended that you keep stripe sizes the same across RAID drive groups.

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 Example of Disk Striping (RAID 0)**



### 2.4.9.1 Stripe Width

Stripe width is the number of drives involved in a drive group where striping is implemented. For example, a four-disk drive group with disk striping has a stripe width of four.

### 2.4.9.2 Stripe Size

The stripe size is the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of disk space and has 16 KB of data residing on each disk in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB.

### 2.4.9.3 Strip Size

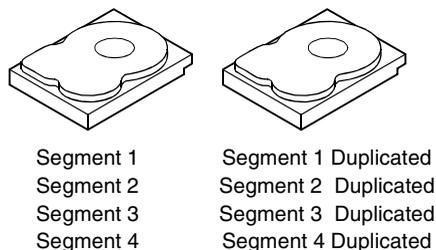
The strip size is the portion of a stripe that resides on a single drive.

## 2.4.10 Disk Mirroring

With mirroring (used in RAID 1 and RAID 10), data written to one drive is simultaneously written to another drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the disk are completely written to a second disk, data is not lost if one disk fails. In addition, both drives contain the same data at all times, so either disk can act as the operational disk. If one disk fails, the contents of the other disk can be used to run the system and reconstruct the failed disk.

Disk mirroring provides 100 percent 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)**



## 2.4.11 Parity

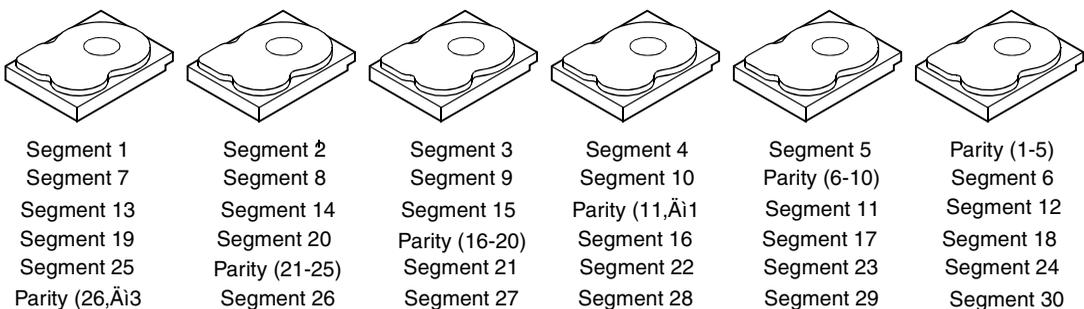
Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent data sets, but parity generation can slow the write process. In RAID, this method is applied to entire drives or stripes across all of the drives in a drive group. The types of parity are described in [Table 2.1](#).

**Table 2.1 Types of Parity**

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

RAID 5 combines distributed parity with disk striping. If a single drive fails, it can be rebuilt from the parity and the data on the remaining drives. An example of a RAID 5 drive group is shown in [Figure 2.3](#). RAID 5 uses parity to provide redundancy for one drive failure without duplicating the contents of entire drives. RAID 6 uses distributed parity and disk striping, also, but adds a second set of parity data so that it can survive up to two drive failures.

**Figure 2.3 Example of Distributed Parity (RAID 5)**



Note: Parity is distributed across all drives in the drive group.

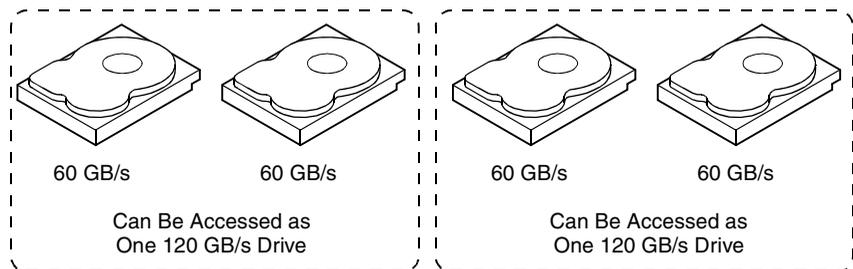
## 2.4.12 Disk Spanning

Disk spanning allows multiple 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 virtual drives must have the same stripe size and must be contiguous. In [Figure 2.4](#), RAID 1 drive groups are turned into a RAID 10 drive group.

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

**Figure 2.4 Example of Disk Spanning**



**Note:** Spanning two contiguous RAID 0 virtual drives does not produce a new RAID level or add fault tolerance. It does increase the capacity of the virtual drive and improves performance by doubling the number of spindles.

### 2.4.12.1 Spanning for RAID 00, RAID 10, RAID 50, and RAID 60

[Table 2.2](#) describes how to use spanning to configure RAID 10, RAID 50, and RAID 60 virtual drives. The virtual drives must have the same stripe size and the maximum number of spans is eight. The full drive capacity is used when you span virtual drives; you cannot specify a smaller drive capacity.

**Table 2.2 Spanning for RAID 10, RAID 50, and RAID 60**

Level	Description
00	Configure RAID 00 by spanning two contiguous RAID 0 virtual drives, up to the maximum number of supported devices for the controller.
10	Configure RAID 10 by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. RAID 10 supports a maximum of eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 virtual drives. The RAID 5 virtual drives must have the same stripe size.
60	Configure RAID 60 by spanning two contiguous RAID 6 virtual drives. The RAID 6 virtual drives must have the same stripe size.

### 2.4.13 Hot Spares

A hot spare is an extra, unused 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. MegaRAID SAS RAID controllers can 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 virtual drives on the same set of drives (a sliced configuration), a rebuild to a hot spare will not occur after a drive failure until the RAID 0 virtual drive is deleted.

The RAID management software allows you to specify 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 virtual drive and marked ready awaiting removal once the rebuild to a hot spare begins. You can make hot spares of the drives that are not in a RAID virtual drive.

You can use the RAID management software to designate the hot spare to have enclosure affinity, meaning that if there are drive failures present on a split backplane configuration, the hot spare will be used first on the backplane side that it resides in.

If the hot spare is designated as having enclosure affinity, it will attempt to rebuild any failed drives on the backplane that it resides in before rebuilding any other drives on other backplanes.

**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:

- Global hot spare
- Dedicated hot spare

#### **2.4.13.1 Global Hot Spare**

A global hot spare drive can be used to replace any failed drive in a redundant drive group 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.

#### **2.4.13.2 Dedicated Hot Spare**

A dedicated hot spare can be used to replace a failed drive only in a selected drive group. One or more drives can be designated as a 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 drive group) are polled every 60 seconds at a minimum, and their status made available in the drive group management software. RAID controllers 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:

- Hot spares are used only in drive groups with redundancy: RAID levels 1, 5, 6, 10, 50, and 60.
- A hot spare connected to a specific RAID controller can be used to rebuild a drive that is connected to the same controller only.

- You must assign the hot spare to one or more drives through the controller BIOS or use drive group management software to place it in the hot spare pool.
- A hot spare must have free space equal to or greater than the drive it replaces. For example, to replace an 18 GB drive, the hot spare must be 18 GB or larger.

## 2.4.14 Disk Rebuilds

When a drive in a RAID drive group fails, you can rebuild the drive by recreating the data that was stored on the drive before it failed. The RAID controller recreates the data using the data stored on the other drives in the drive group. Rebuilding can be done only in drive groups with data redundancy, which includes RAID 1, 5, 6, 10, 50, and 60 drive groups.

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 that the data on the failed drive can be rebuilt.

The failed drive is removed from the virtual drive and marked ready awaiting removal when 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 virtual drive before management applications 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 source drive fails during a rebuild to a hot spare, the rebuild fails, and the failed source drive is marked as offline. In addition, the rebuilding hot spare drive is changed back to a hot spare. After a rebuild fails because of a source drive failure, the dedicated hot spare is still dedicated and assigned to the correct drive group, and the global hot spare is still global.

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

### 2.4.15 Rebuild Rate

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

### 2.4.16 Hot Swap

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

- The newly inserted drive is the same capacity as or larger than the failed drive
- 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 drives and rebuild the contents of the drive automatically.

### 2.4.17 Drive States

A drive state is a property indicating the status of the drive. The drive states are described in [Table 2.3](#).

**Table 2.3 Drive States**

State	Description
Online	A drive that can be accessed by the RAID controller and is part of the virtual drive.
Unconfigured Good	A drive that is functioning normally but is not configured as a part of a virtual drive or as a hot spare.

**Table 2.3 Drive States (Cont.)**

State	Description
Hot Spare	A drive that is powered up and ready for use as a spare in case an online drive fails.
Failed	A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.
Unconfigured Bad	A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
Missing	A drive that was Online but which has been removed from its location.
Offline	A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned.  Note: When a virtual drive with cached data goes offline, the cache for the virtual drive is discarded. Because the virtual drive is offline, the cache cannot be saved.

## 2.4.18 Virtual Drive States

The virtual drive states are described in [Table 2.4](#).

**Table 2.4 Virtual Drive States**

State	Description
Optimal	The virtual drive operating condition is good. All configured drives are online.
Degraded	The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
Partial Degraded	The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.

## 2.4.19 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.

---

## 2.5 RAID Levels

The RAID controller supports RAID levels 0, 00, 1, 5, 6, 10, 50, and 60. The supported RAID levels are summarized in the following section. In addition, it supports independent drives (configured as RAID 0 and RAID 00.) The following sections describe the RAID levels in detail.

### 2.5.1 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.

A RAID 00 drive group is a spanned drive group that creates a striped set from a series of RAID 0 drive groups.

RAID 1 uses mirroring so that data written to one drive is simultaneously written to another 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 6 uses distributed parity, with two independent parity blocks per stripe, and disk striping. A RAID 6 virtual drive can survive the loss of two drives without losing data. A RAID 6 drive group, which requires a minimum of three drives, is similar to a RAID 5 drive group. Blocks of data and parity information are written across all drives. The parity information is used to recover the data if one or two drives fail in the drive group.

RAID 10, a combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. A RAID 10 drive group is a spanned drive group that creates a striped set from a series of mirrored drives. RAID 10 allows a maximum of eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. RAID 10 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. A RAID 50 drive group is a spanned drive group in which data is striped across multiple RAID 5 drive groups. RAID 50 works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

**Note:** Having virtual drives of different RAID levels, such as RAID 0 and RAID 5, in the same drive group is not allowed. For example, if an existing RAID 5 virtual drive is created out of partial space in an array, the next virtual drive in the array has to be R5 only.

RAID 60, a combination of RAID 0 and RAID 6, uses distributed parity, with two independent parity blocks per stripe in each RAID set, and disk striping. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. It works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

## 2.5.2 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 drive group depends on a number of factors:

- The number of drives in the drive group
- The capacity of the drives in the drive group
- The need for data redundancy
- The disk performance requirements

## 2.5.3 RAID 0

RAID 0 provides disk striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy, but, along with RAID 0, does offer the best performance of any RAID level. RAID 0 breaks up data into smaller segments, and then stripes the data segments across each drive in the drive group. The size of each data segment is determined by the stripe size. RAID 0 offers high bandwidth.

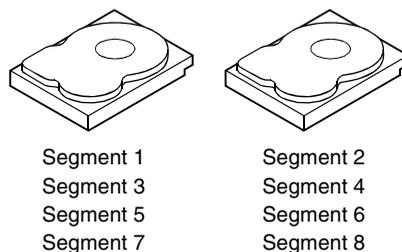
**Note:** RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) will fail.

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA 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. [Table 2.5](#) provides an overview of RAID 0. [Figure 2.5](#) provides a graphic example of a RAID 0 drive group.

**Table 2.5 RAID 0 Overview**

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

**Figure 2.5 RAID 0 Drive Group Example with Two Drives**



## 2.5.4 RAID 00

A RAID 00 drive group is a spanned drive group that creates a striped set from a series of RAID 0 drive groups. RAID 00 does not provide any data redundancy, but, along with RAID 0, does offer the best performance of any RAID level. RAID 00 breaks up data into smaller segments and then stripes the data segments across each drive in the drive groups. The size of each data segment is determined by the stripe size. RAID 00 offers high bandwidth.

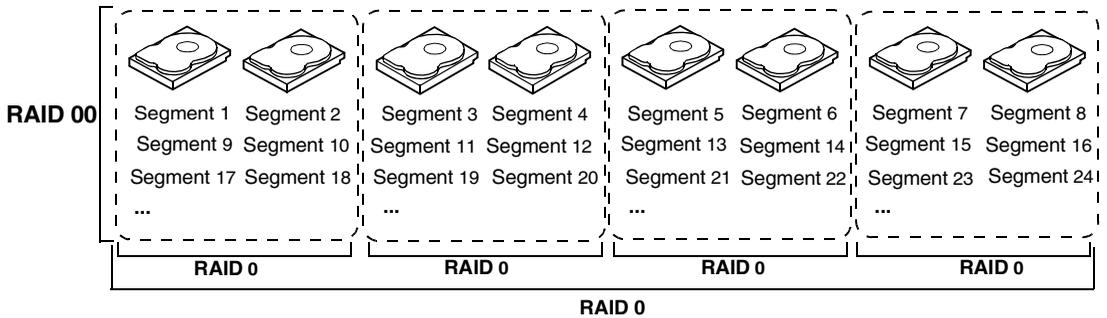
**Note:** RAID level 00 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) will fail.

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. RAID 00 involves no parity calculations to complicate the write operation. This makes RAID 00 ideal for applications that require high bandwidth but do not require fault tolerance. [Table 2.6](#) provides an overview of RAID 00. [Figure 2.6](#) provides a graphic example of a RAID 00 drive group.

**Table 2.6 RAID 00 Overview**

<b>Uses</b>	Provides high data throughput, especially for large files. Any environment that does not require fault tolerance.
<b>Strong Points</b>	Provides increased data throughput for large files. No capacity loss penalty for parity.
<b>Weak Points</b>	Does not provide fault tolerance or high bandwidth. All data lost if any drive fails.
<b>Drives</b>	2 to 256

**Figure 2.6 RAID 0 Drive Group Example with Two Drives**



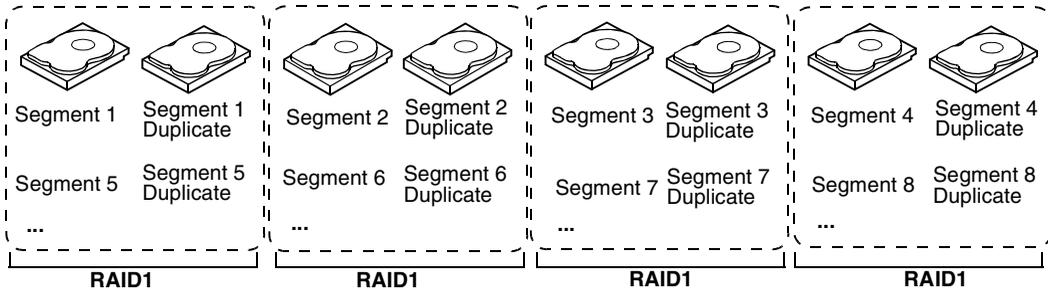
## 2.5.5 RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive in the drive group. RAID 1 supports an even number of drives from 2 to 32 in a single span. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. [Table 2.7](#) provides an overview of RAID 1. [Figure 2.7](#) provides a graphic example of a RAID 1 drive group.

**Table 2.7 RAID 1 Overview**

<b>Uses</b>	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
<b>Strong Points</b>	Provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
<b>Weak Points</b>	Requires twice as many drives. Performance is impaired during drive rebuilds.
<b>Drives</b>	2 - 32 (must be an even number of drives)

**Figure 2.7 RAID 1 Drive Group**



## 2.5.6 RAID 5

RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking is used to detect errors in the data. In RAID 5, the parity information is written to all 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.

[Table 2.8](#) provides an overview of RAID 5. [Figure 2.8](#) provides a graphic example of a RAID 5 drive group.

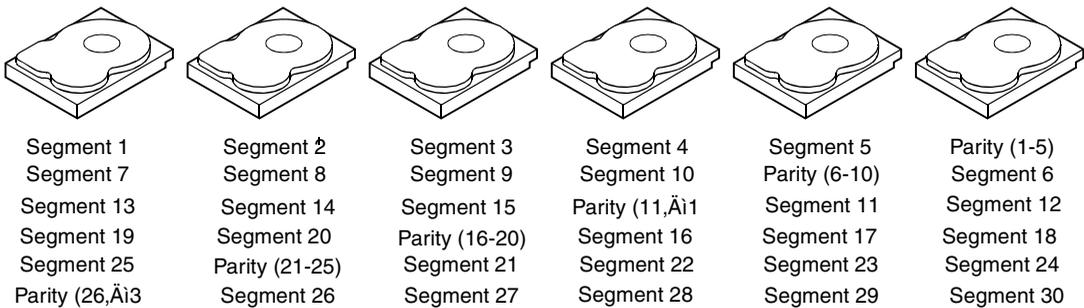
**Table 2.8 RAID 5 Overview**

<b>Uses</b>	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.
<b>Strong Points</b>	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.

**Table 2.8 RAID 5 Overview**

	Not well-suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). 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.
<b>Weak Points</b>	
<b>Drives</b>	3 to 32

**Figure 2.8 RAID 5 Drive Group with Six Drives**



Note: Parity is distributed across all drives in the drive group.

## 2.5.7 RAID 6

RAID 6 is similar to RAID 5 (disk striping and parity), except that instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two drives in a virtual drive without losing data. Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 6 for data that requires a very high level of protection from loss.

In the case of a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive, and then the other failed drive.

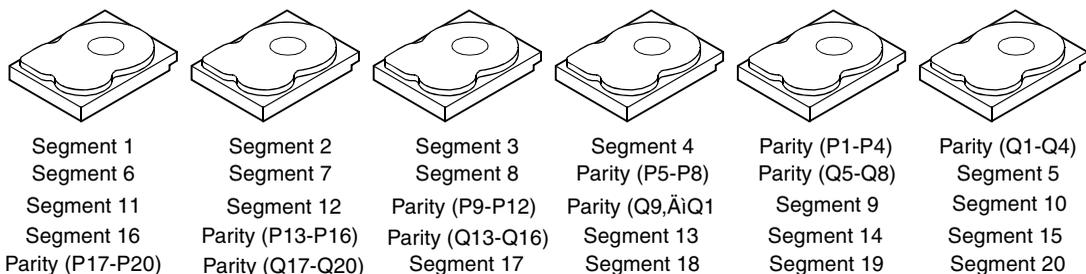
[Table 2.7](#) provides a graphic example of a RAID 6 drive group.

**Table 2.9 RAID 6 Overview**

<b>Uses</b>	Use 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.
<b>Strong Points</b>	Provides data redundancy, high read rates, and good performance in most environments. Can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 5.
<b>Weak Points</b>	Not well-suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.
<b>Drives</b>	3 to 32

Figure 2.9 shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.

**Figure 2.9 Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)**



Note: Parity is distributed across all drives in the drive group.

## 2.5.8 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1, and consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks and then mirrors the blocks of data to each RAID 1 drive group. The first

RAID 1 drive in each drive group then duplicates its data to the second drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. The RAID 1 virtual drives must have the same stripe size.

Spanning is used because one virtual drive is defined across more than one drive group. Virtual drives defined across multiple RAID 1 level drive groups are referred to as RAID level 10, (1+0). Data is striped across drive groups to increase performance by enabling access to multiple drive groups simultaneously.

Each spanned RAID 10 virtual drive can tolerate multiple drive failures, as long as each failure is in a separate drive group. If there are drive failures, less than total drive capacity is available.

Configure RAID 10 by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. RAID 10 supports a maximum of eight spans, with a maximum of 32 drives per span. You must use an even number of drives in each RAID 10 virtual drive in the span.

**Note:** Other factors, such as the type of controller, can restrict the number of drives supported by RAID 10 virtual drives.

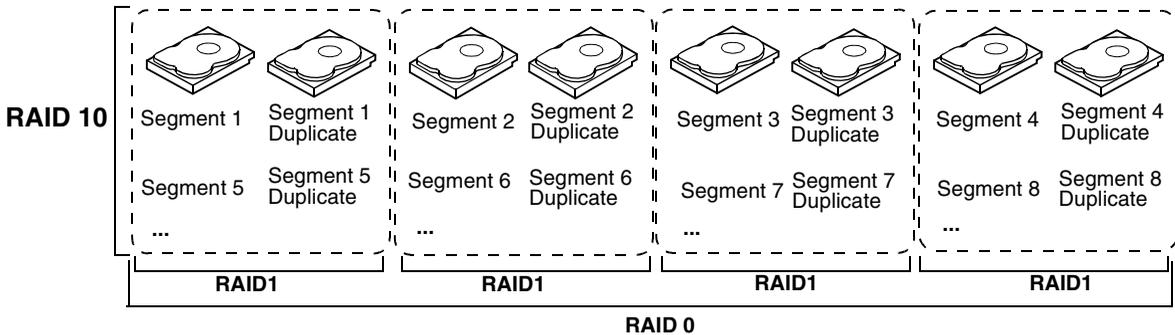
Table 2.10 provides an overview of RAID 10.

**Table 2.10 RAID 10 Overview**

<b>Uses</b>	Appropriate when used with data storage that needs 100 percent redundancy of mirrored drive groups and that also needs the enhanced I/O performance of RAID 0 (striped drive groups.) RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.
<b>Strong Points</b>	Provides both high data transfer rates and complete data redundancy.
<b>Weak Points</b>	Requires twice as many drives as all other RAID levels except RAID 1.
<b>Drives</b>	4 - the maximum number of drives supported by the controller (with a maximum of eight spans)

In Figure 2.10, virtual drive 0 is created by distributing data across four drive groups (drive groups 0 through 3).

**Figure 2.10 RAID 10 Level Virtual Drive**



## 2.5.9 RAID 50

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

RAID 50 breaks up data into smaller blocks and 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 and then writes the blocks of data and parity to each drive in the drive group. 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 eight spans and tolerate up to eight drive failures, though less than total drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

[Table 2.11](#) provides an overview of RAID 50.

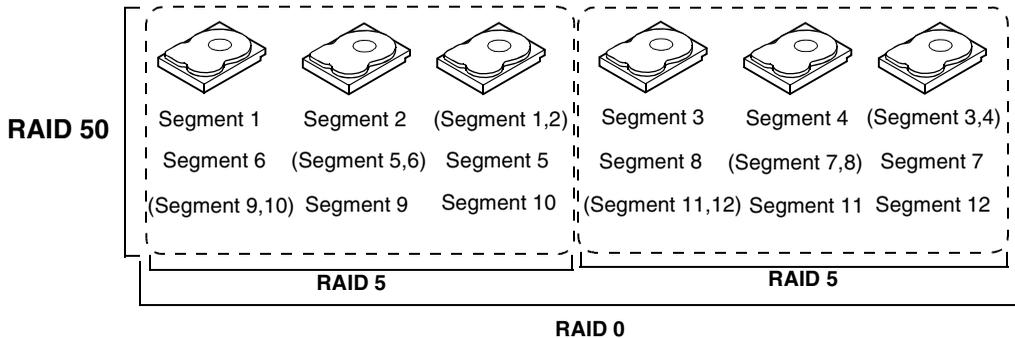
**Table 2.11 RAID 50 Overview**

<b>Uses</b>	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.
<b>Strong Points</b>	Provides high data throughput, data redundancy, and very good performance.
<b>Weak Points</b>	Requires 2 to 8 times as many parity drives as RAID 5.

**Table 2.11 RAID 50 Overview**

<b>Drives</b>	Eight spans of RAID 5 drive groups containing 3-32 drives each (limited by the maximum number of devices supported by the controller)
---------------	---

**Figure 2.11 RAID 50 Level Virtual Drive**



## 2.5.10 RAID 60

RAID 60 provides the features of both RAID 0 and RAID 6, and includes both parity and disk striping across multiple drive groups. RAID 6 supports two independent parity blocks per stripe. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. RAID 60 is best implemented on two RAID 6 drive groups with data striped across both drive groups.

RAID 60 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 6 disk set. RAID 6 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

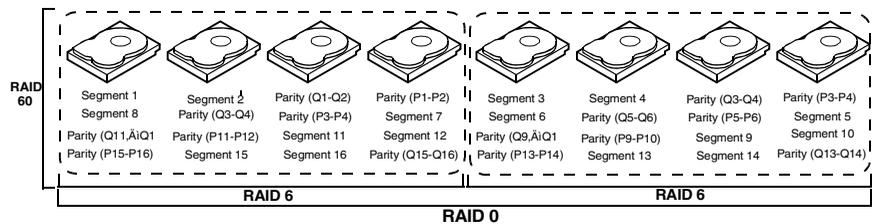
RAID 60 can support up to 8 spans and tolerate up to 16 drive failures, though less than total drive capacity is available. Two drive failures can be tolerated in each RAID 6 level drive group.

**Table 2.12 RAID 60 Overview**

<b>Uses</b>	<p>Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 60 for data that requires a very high level of protection from loss.</p> <p>In the case of a failure of one drive or two drives in a RAID set in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 set in a RAID 60 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds can occur at the same time.</p> <p>Use 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.</p>
<b>Strong Points</b>	<p>Provides data redundancy, high read rates, and good performance in most environments. Each RAID 6 set can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 50, though random reads in RAID 60 might be slightly faster because data is spread across at least one more disk in each RAID 6 set.</p>
<b>Weak Points</b>	<p>Not well suited to tasks requiring lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.</p>
<b>Drives</b>	<p>A minimum of 8</p>

Figure 2.12 shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.

**Figure 2.12 RAID 60 Level Virtual Drive**



Note: Parity is distributed across all drives in the drive group.

## 2.6 RAID Configuration Strategies

The most important factors in RAID drive group configuration are:

- Virtual drive availability (fault tolerance)
- Virtual drive performance
- Virtual drive capacity

You cannot configure a virtual drive that optimizes all three factors, but it is easy to choose a virtual 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 virtual drive availability (fault tolerance), virtual drive performance, and virtual drive capacity.

### 2.6.1 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 the RAID controller instantly plugs into the system when an active drive fails. After the hot spare is automatically moved into the RAID drive group, the failed drive is automatically rebuilt on the spare drive. The RAID drive group 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

WebBIOS Configuration Utility allows a failed drive to be replaced and automatically rebuilt by “hot-swapping” the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs, providing a high degree of fault tolerance and zero downtime.

**Table 2.13 RAID Levels and Fault Tolerance**

<b>RAID Level</b>	<b>Fault Tolerance</b>
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance.
00	Does not provide fault tolerance. All data in a virtual drive is lost if any drive in that virtual drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 00 is ideal for applications that require high bandwidth but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Since the contents of the 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 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 drives in a drive group. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.
6	Combines distributed parity with disk striping. RAID 6 can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.
10	Provides complete data redundancy using striping across spanned RAID 1 drive groups. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. RAID 10 can sustain a drive failure in each mirrored drive group and maintain drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. 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 drive group and still maintain data integrity.

RAID Level	Fault Tolerance
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. RAID 60 can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. RAID 60 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.

## 2.6.2 Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O is faster because drives can be accessed simultaneously. [Table 2.14](#) describes the performance for each RAID level.

**Table 2.14 RAID Levels and Performance**

RAID Level	Performance
0	RAID 0 (striping) offers excellent performance. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the drive group. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
00	RAID 00 (striping in a spanned drive group) offers excellent performance. RAID 00 breaks up data into smaller blocks and then writes a block to each drive in the drive groups. Disk striping writes data across multiple drives instead of just one drive. Striping involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 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.

RAID Level	Performance
5	<p>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.</p> <p>Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. 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.</p>
6	<p>RAID 6 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. However, RAID 6 is not well suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>
10	<p>RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive 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 drive group.</p>
50	<p>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 capacity of the virtual drive 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 drive group.</p>

RAID Level	Performance
60	<p>RAID 60 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 capacity of the virtual drive 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 6 drive group.</p> <p>RAID 60 is not well suited to tasks requiring a lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>

### 2.6.3 Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. There are several variables to consider. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity (RAID 5 or RAID 6). RAID 5, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than RAID 1. [Table 2.15](#) explains the effects of the RAID levels on storage capacity.

**Table 2.15 RAID Levels and Capacity**

RAID Level	Capacity
0	RAID 0 (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 drives.
00	RAID 00 (striping in a spanned drive group) 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 00 provides maximum storage capacity for a given set of drives.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.

RAID Level	Capacity
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire 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 drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 6 more expensive to implement.
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 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.
60	RAID 60 provides redundancy for two drive failures in each RAID set without duplicating the contents of entire drives. However, it requires extra capacity because a RAID 60 virtual drive has to generate two sets of parity data for each write operation. This makes RAID 60 more expensive to implement.

---

## 2.7 RAID Availability

### 2.7.1 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 failed 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 drive problems, while keeping the servers running and data available. The following subsections describe these features.

### 2.7.1.1 Spare Drives

You can use spare drives to replace failed or defective drives in a drive group. A replacement drive must be at least as large as the 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 drives that power up along with the RAID drives and operate in a standby state. If a drive used in a RAID virtual 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, 6, 10, 50, and 60.

**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."

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

### 2.7.1.2 Rebuilding

If a drive fails in a drive group that is configured as a RAID 1, 5, 6, 10, 50, or 60 virtual 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 drives. 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.

---

## 2.8 Configuration Planning

Factors to consider when planning a configuration are the number of drives the RAID controller can support, the purpose of the drive group, 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.

### 2.8.1 Number of Drives

Your configuration planning for the SAS RAID controller depends in part on the number of drives that you want to use in a RAID drive group. The number of drives in a drive group determines the RAID levels that can be supported. Only one RAID level can be assigned to each virtual drive.

### 2.8.2 Drive Group Purpose

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

- Will this drive group increase the system storage capacity for general-purpose file and print servers? Use RAID 5, 6, 10, 50, or 60.
- Does this drive group support any software system that must be available 24 hours per day? Use RAID 1, 5, 6, 10, 50, or 60.
- Will the information stored in this drive group contain large audio or video files that must be available on demand? Use RAID 0 or 00.
- Will this drive group contain data from an imaging system? Use RAID 0, 00, or 10.

Fill out [Table 2.16](#) to help you plan the drive group configuration. Rank the requirements for your drive group, such as storage space and data redundancy, in order of importance, and then review the suggested RAID levels.

**Table 2.16 Factors to Consider for Drive Group Configuration**

<b>Requirement</b>	<b>Rank</b>	<b>Suggested RAID Level(s)</b>
Storage space		RAID 0, RAID 00, RAID 5
Data redundancy		RAID 5, RAID 6, RAID 10, RAID 50, RAID 60
Drive performance and throughput		RAID 0, RAID 00, RAID 10
Hot spares (extra drives required)		RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60

# Chapter 3

## WebBIOS Configuration Utility

---

This chapter describes the WebBIOS Configuration Utility (CU) and consists of the following sections:

- [Section 3.1, “Overview”](#)
  - [Section 3.2, “Starting the WebBIOS CU”](#)
  - [Section 3.3, “WebBIOS CU Main Screen Options”](#)
  - [Section 3.4, “Creating a Storage Configuration”](#)
  - [Section 3.5, “Viewing and Changing Device Properties”](#)
  - [Section 3.6, “Viewing System Event Information”](#)
  - [Section 3.7, “Managing Configurations”](#)
- 

### 3.1 Overview

The WebBIOS CU enables you to create and manage RAID configurations on LSI SAS controllers. Unlike the MegaRAID Storage Manager™ software, the WebBIOS CU resides in the SAS controller BIOS and operates independently of the operating system.

You can use the WebBIOS CU to do the following tasks:

- Create drive groups and virtual drives for storage configurations
- Display controller, virtual drive, drive, and battery backup unit (BBU) properties, and change parameters
- Delete virtual drives
- Migrate a storage configuration to a different RAID level
- Detect configuration mismatches
- Import a foreign configuration

- Scan devices connected to the controller
- Initialize virtual drives
- Check configurations for data consistency

The WebBIOS CU provides a configuration wizard to guide you through the configuration of virtual drives and drive groups.

---

## 3.2 Starting the WebBIOS CU

Follow these steps to start the WebBIOS CU and access the main screen.

1. When the host computer is booting, hold down the <Ctrl> key and press the <H> key when the following text appears on the screen:

```
Copyright© LSI Corporation  
Press <Ctrl><H> for WebBIOS
```

The Controller Selection screen appears.

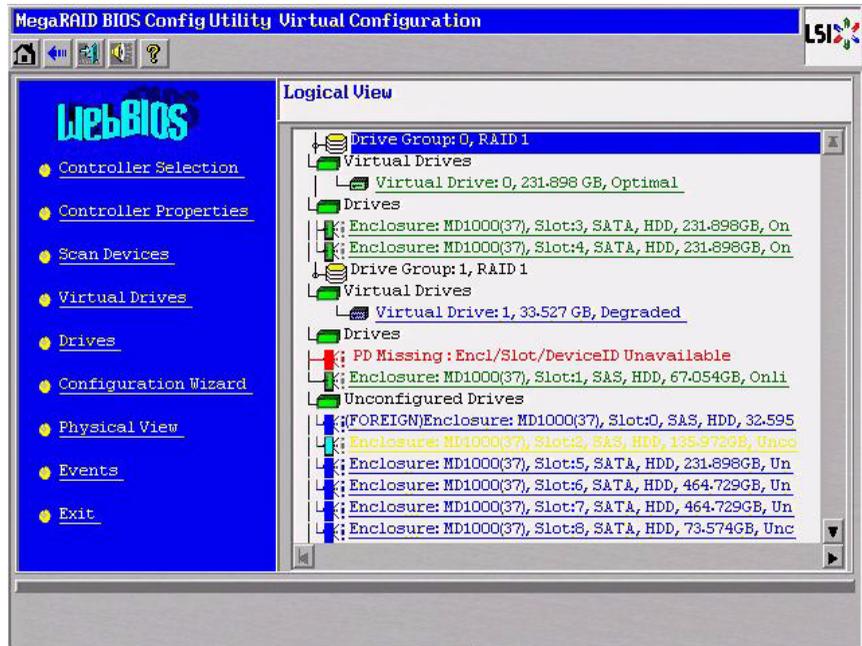
2. If the system has multiple SAS controllers, select a controller.
3. Click **Start** to continue.

The main WebBIOS CU screen appears.

### 3.3 WebBIOS CU Main Screen Options

Figure 3.1 shows the screen that appears when you start the WebBIOS CU and select a controller.

Figure 3.1 WebBIOS CU Main Screen



In the right frame, the screen shows the virtual drives configured on the controller, and the drives that are connected to the controller. In addition, the screen identifies drives that are foreign or missing.

**Note:** In the list of virtual drives, the drive nodes are sorted based on the order in which you added the drives to the drive group, rather than the physical slot order that displays in the physical trees.

**Note:** The minimum screen resolution for WebBIOS is 640x480.

To toggle between the physical view and logical view of the storage devices connected to the controller, click **Physical View** or **Logical View** in the menu on the left. When the physical view screen appears, it shows the drive groups that are configured on this controller.

For drives in an enclosure, the screen shows the following drive information:

- Enclosure
- Slot
- Interface type (such as SAS or SATA)
- Drive type (HDD or SSD)
- Drive size
- Drive status (such as **Online** or **Unconfigured Good**)

The toolbar at the top of the WebBIOS CU has the following buttons, as listed in [Table 3.1](#).

**Table 3.1 WebBIOS CU Toolbar Icons**

Icon	Description
	Click this icon to return to the main screen from any other WebBIOS CU screen.
	Click this icon to return to the previous screen that you were viewing.
	Click this icon to exit the WebBIOS CU program.
	Click this icon to turn off the sound on the onboard controller alarm.
	Click this icon to display information about the WebBIOS CU version, browser version, and HTML interface engine.

Here is a description of the options listed on the left of the main WebBIOS CU screen:

- **Controller Selection:** Select this option to view the Controller Selection screen, where you can select a different controller. You can then view information about the controller and the devices connected to it, or create a new configuration on the controller.
- **Controller Properties:** Select this option to view the properties of the currently selected SAS controller. For more information, see [Section 3.5.1, “Viewing and Changing Controller Properties.”](#)
- **Scan Devices:** Select this option to have the WebBIOS CU re-scan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS CU displays the results of the scan in the physical and virtual drive descriptions.
- **Virtual Drives:** Select this option to view the Virtual Drives screen, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see [Section 3.5.2, “Viewing and Changing Virtual Drive Properties.”](#)
- **Drives:** Select this option to view the Drives screen, where you can view drive properties, create hot spares, and perform other tasks. For more information, see [Section 3.5.3, “Viewing Drive Properties.”](#)
- **Configuration Wizard:** Select this option to start the Configuration Wizard and create a new storage configuration, clear a configuration, or add a configuration. For more information, see [Section 3.4, “Creating a Storage Configuration.”](#)
- **Physical View/Logical View:** Select this option to toggle between the Physical View screen and the Logical View screen.
- **Events:** Select this option to view system events in the Event Information screen. For more information, see [Section 3.6, “Viewing System Event Information.”](#)
- **Exit:** Select this option to exit the WebBIOS CU and continue with system boot.

---

## 3.4 Creating a Storage Configuration

This section explains how to use the WebBIOS CU Configuration Wizard to configure RAID drive groups and virtual drives. The following subsections explain how to use the Configuration Wizard to create storage configurations:

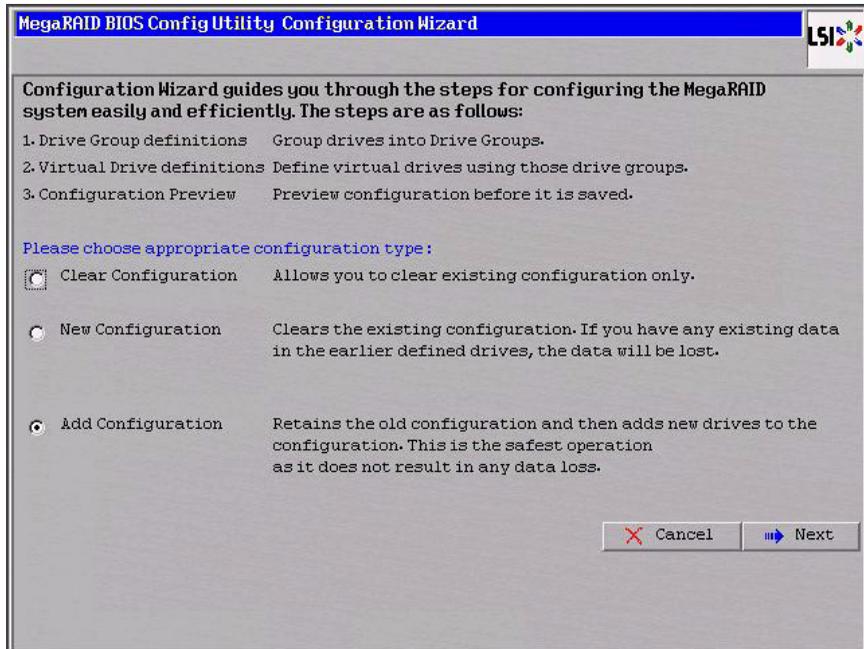
- [Section 3.4.1, “Selecting the Configuration with the Configuration Wizard”](#)
- [Section 3.4.2, “Using Automatic Configuration”](#)
- [Section 3.4.3, “Using Manual Configuration”](#)

### 3.4.1 Selecting the Configuration with the Configuration Wizard

Follow these steps to start the Configuration Wizard, and select a configuration option and mode:

1. Click **Configuration Wizard** on the WebBIOS main screen.  
The first Configuration Wizard screen appears, as shown in [Figure 3.2](#).

**Figure 3.2 WebBIOS Configuration Wizard Screen**



2. Select a configuration option.

**Caution:** If you choose the first or second option, all existing data in the configuration will be deleted. Make a backup of any data that you want to keep before you choose an option.

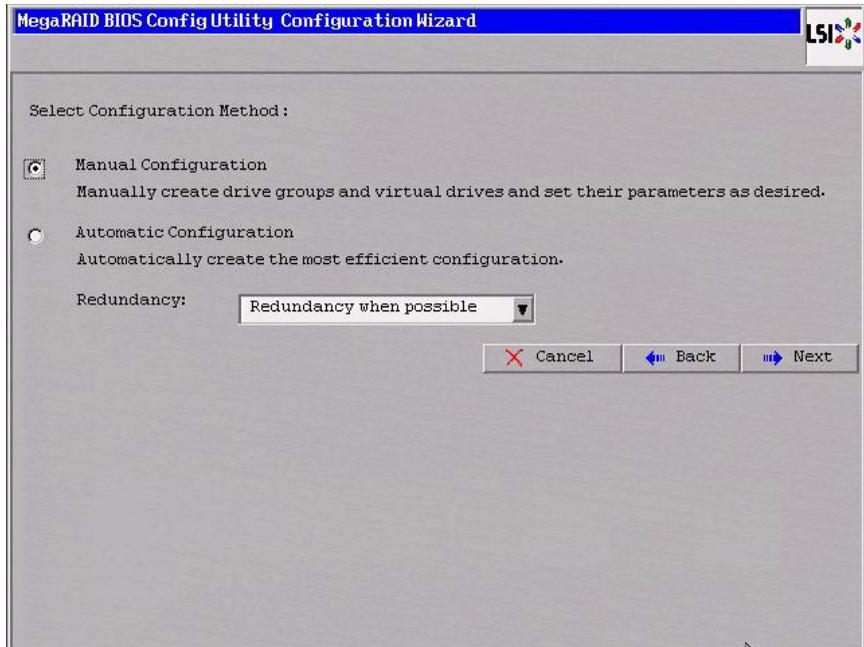
- **Clear Configuration:** Clears the existing configuration.
- **New Configuration:** Clears the existing configuration and lets you create a new configuration.
- **Add Configuration:** Retains the existing storage configuration and adds new drives to it (this does not cause any data loss).

3. Click **Next**.

A dialog box warns that you will lose data if you select Clear Configuration or New Configuration.

The WebBIOS Configuration Method screen appears, as shown in [Figure 3.3](#).

**Figure 3.3 WebBIOS Configuration Method Screen**



4. On this screen, select a configuration mode:
  - **Manual Configuration:** Allows you to control all attributes of the new storage configuration as you create drive groups and virtual drives, and set their parameters.
  - **Automatic Configuration:** Automatically creates an optimal RAID configuration.
5. If you select Automatic Configuration, you can choose whether to create a redundant RAID drive group or a non-redundant RAID 0 drive group. Select one of the following options in the Redundancy field:
  - **Redundancy when possible**
  - **No redundancy**
6. Click **Next** to continue.

If you select the Automatic Configuration option, continue with [Section 3.4.2, “Using Automatic Configuration.”](#) If you select Manual Configuration, continue with [Section 3.4.3, “Using Manual Configuration.”](#)

## 3.4.2 Using Automatic Configuration

Follow these instructions to create a configuration with automatic configuration, either with or without redundancy:

1. When WebBIOS shows the proposed new configuration, review the information on the screen, and click **Accept** to accept it. (Or click **Back** to go back and change the configuration.)
  - **RAID 0:** If you select **Automatic Configuration** and **No Redundancy**, WebBIOS creates a RAID 0 configuration.
  - **RAID 1:** If you select **Automatic Configuration** and **Redundancy when possible**, and only two drives are available, WebBIOS creates a RAID 1 configuration.
  - **RAID 5:** If you select **Automatic Configuration** and **Redundancy when possible**, and three or more drives are available, WebBIOS creates a RAID 5 configuration.
  - **RAID 6:** If you select **Automatic Configuration** and **Redundancy when possible**, and the RAID 6 option is enabled, and three or more drives are available, WebBIOS creates a RAID 6 configuration.
2. Click **Yes** when you are prompted to save the configuration.
3. Click **Yes** when you are prompted to initialize the new virtual drive(s).  
WebBIOS CU begins a background initialization of the virtual drives.

## 3.4.3 Using Manual Configuration

The following subsections contain the procedures for creating RAID drive groups for RAID levels 0, 00, 1, 5, 6, 10, 50, and 60:

- [Section 3.4.3.1, “Using Manual Configuration: RAID 0”](#)
- [Section 3.4.3.2, “Using Manual Configuration: RAID 00”](#)
- [Section 3.4.3.3, “Using Manual Configuration: RAID 1”](#)
- [Section 3.4.3.4, “Using Manual Configuration: RAID 5”](#)
- [Section 3.4.3.5, “Using Manual Configuration: RAID 6”](#)
- [Section 3.4.3.6, “Using Manual Configuration: RAID 10”](#)
- [Section 3.4.3.7, “Using Manual Configuration: RAID 50”](#)

- [Section 3.4.3.8, “Using Manual Configuration: RAID 60”](#)

### 3.4.3.1 Using Manual Configuration: RAID 0

RAID 0 provides drive striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy but does offer excellent performance. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance. RAID 0 also denotes an independent or single drive.

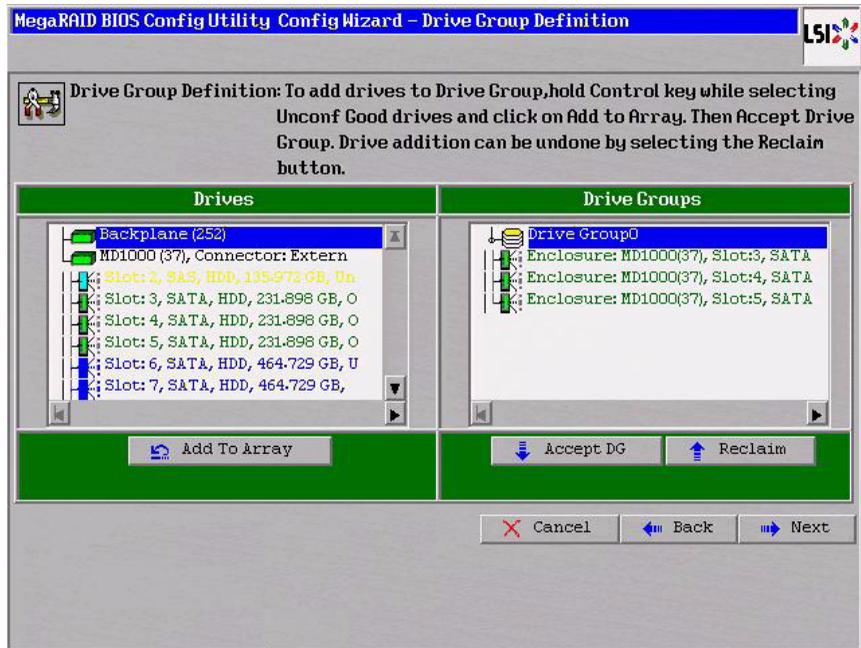
**Note:** RAID level 0 is not fault-tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) fails.

When you select **Manual Configuration** and click **Next**, the drive group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting two or more ready drives in the Drives panel on the left until you have selected all desired drives for the drive group.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in [Figure 3.4](#).

If you need to undo the changes, click the **Reclaim** button.

Figure 3.4 WebBIOS Drive Group Definition Screen



3. When you have finished selecting drives for the drive group, click **Accept DG**.
4. Click **Next**.

The Virtual Drive Definition screen appears, as shown in [Figure 3.5](#). This screen lists the possible RAID levels for the drive group. Use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

**Figure 3.5 WebBIOS Virtual Drive Definition Screen**

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level	RAID 0
Strip Size	64 KB
Access Policy	RW
Read Policy	Normal
Write Policy	WThru
<input checked="" type="checkbox"/> Wthru for BAD BBU	
IO Policy	Direct
Drive Cache	NoChange
Disable BGI	No
Select Size	237464 MB

Press Back Button To Add Another Virtual Drive.

Reclaim

Cancel Back Next

5. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 0.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

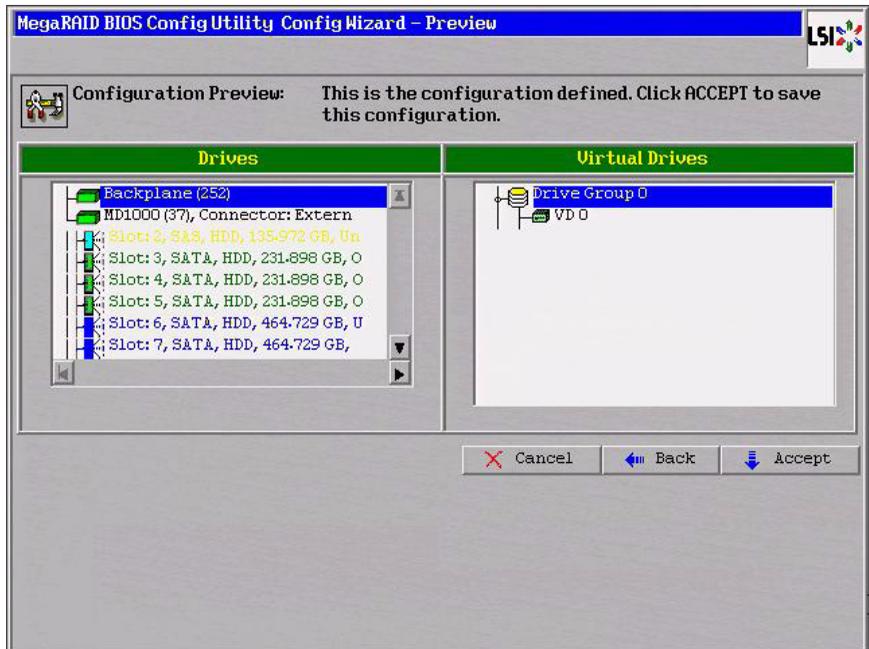
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access. This is the default.
  - ◇ *Read Only:* Allow read-only access.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In cached I/O mode, all reads are buffered in cache memory.
  - **Drive Cache:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 0 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
6. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
  7. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.6](#).

**Figure 3.6 RAID 0 Configuration Preview**



8. Check the information in the configuration preview.
9. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
10. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

### 3.4.3.2 Using Manual Configuration: RAID 00

A RAID 00 drive group is a spanned drive group that creates a striped set from a series of RAID 0 drive groups. It breaks up data into smaller blocks and then stripes the blocks of data to RAID 00 drive groups. The size of each block is determined by the stripe size parameter, which is 64 Kbytes.

RAID 00 does not provide any data redundancy but does offer excellent performance. RAID 00 is ideal for applications that require high bandwidth but do not require fault tolerance.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears.

You use the Disk Group Definition screen to select drives to create drive groups.

1. Hold <Ctrl> while you select ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

3. Click **Accept DG** to create a RAID 0 drive group.

An icon for the next drive group appears in the right panel.

4. Hold <Ctrl> while you select more ready drives in the Drives panel to create a second RAID 0 drive group.

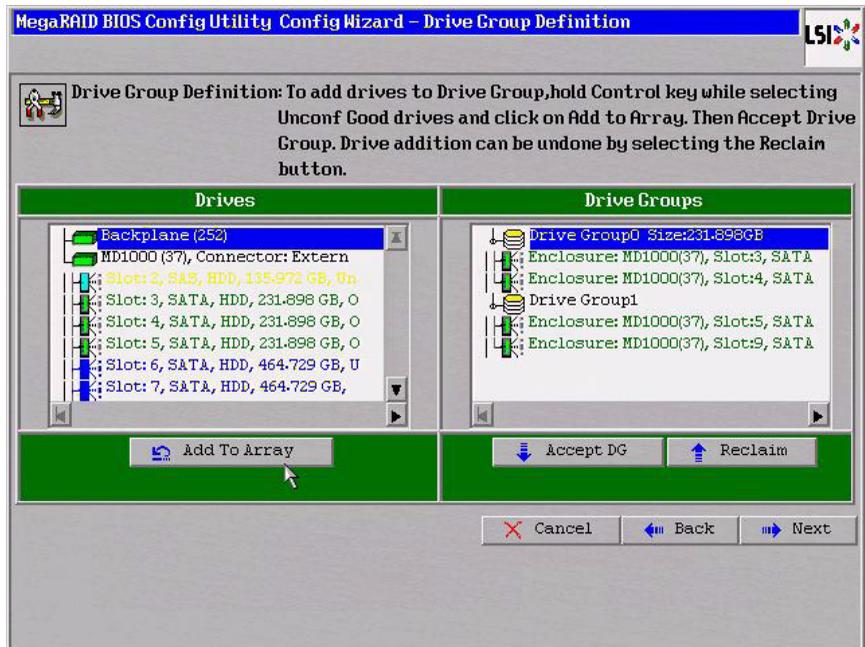
5. Click **Add To Array** to move the drives to a second drive group configuration in the Disk Groups panel, as shown in [Figure 3.20](#).

If you need to undo the changes, click the **Reclaim** button.

**Note:** RAID 00 supports a maximum of eight spans, with a maximum of 32 drives per span. (Other factors, such as the type of controller, can limit the number of drives.)

6. Click **Accept DG** to create a RAID 0 drive group.

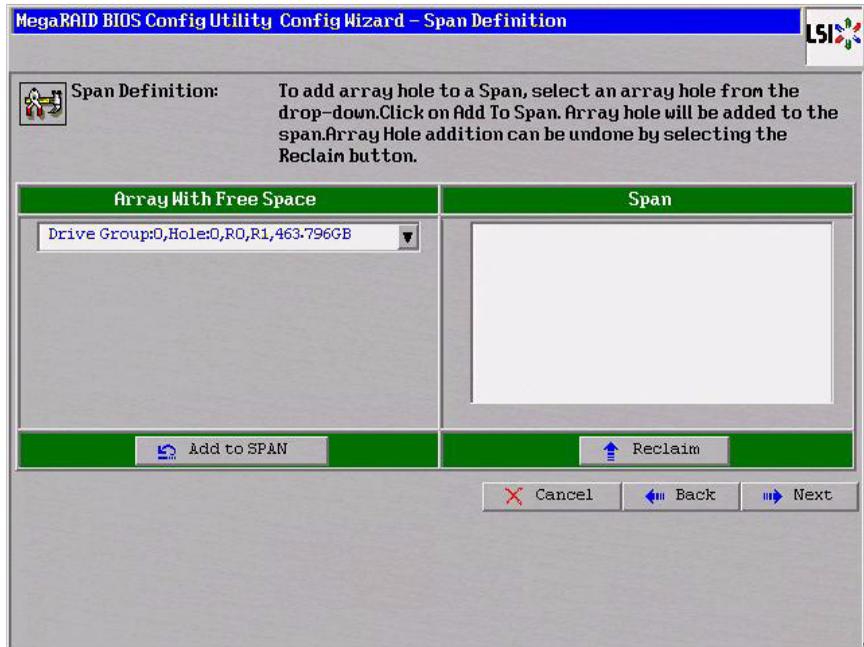
Figure 3.7 WebBIOS Drive Group Definition Screen



7. Repeat [step 4](#) through [step 6](#) until you have selected all the drives you want for the drive groups.
8. When you have finished selecting drives for the drive groups, select each drive group and then click **Accept DG** for each selection.
9. Click **Next**.

The Span Definition screen appears, as shown in [Figure 3.8](#). This screen shows the drive group holes that you can select to add to a span.

**Figure 3.8 WebBIOS Span Definition Screen**



10. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group, and then click **Add to SPAN**.

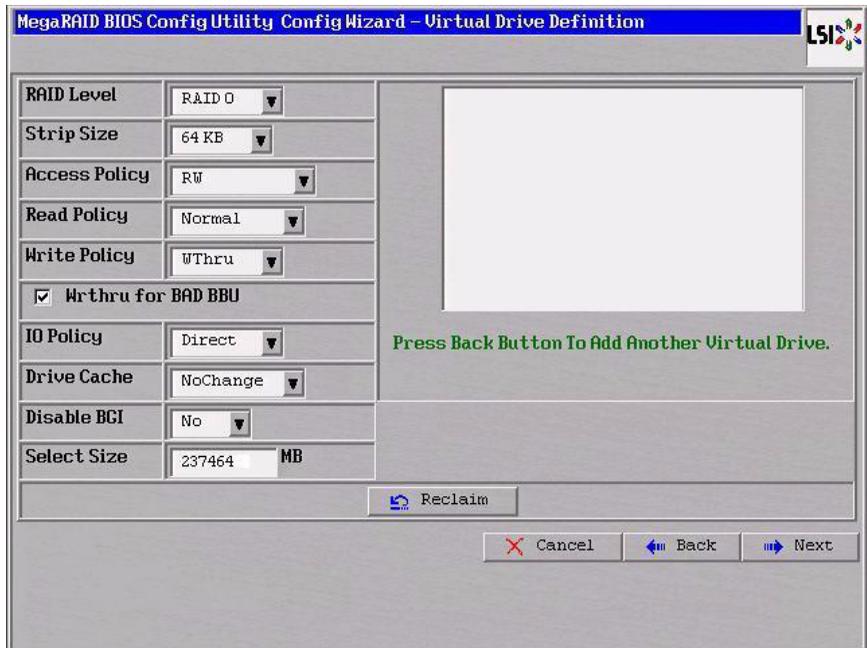
The drive group you select displays in the right frame under the heading **Span**.

11. Hold <Ctrl> while you select a second drive group, and then click **Add to SPAN**.
12. Repeat [step 10](#) until you have selected all of the drive groups that you want.
13. Click **Next**.

The Virtual Drive Definition screen appears, as shown in [Figure 3.22](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

14. Hold <Ctrl> while you select drive groups in the Configuration panel on the right.

**Figure 3.9 WebBIOS Virtual Drive Definition Screen**



15. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 0.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

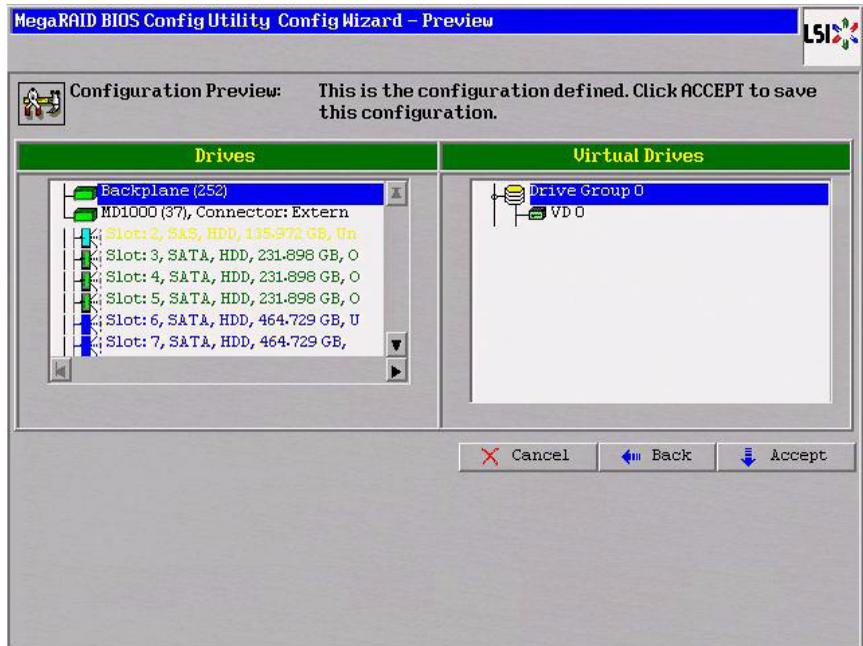
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access.
  - ◇ *Read Only:* Allow read-only access. This type of access is the default.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This option disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This option enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. The policy does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, the block comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Cache:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 00 shown in the Configuration Panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
16. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
17. When you are finished defining virtual drives, click **Next**.

The Configuration Preview screen appears, as shown in [Figure 3.23](#).

**Figure 3.10 RAID 00 Configuration Preview**



18. Check the information in the configuration preview.
19. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation and return to the WebBIOS main menu, or click **Back** to return to the previous screens and change the configuration.
20. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

### 3.4.3.3 Using Manual Configuration: 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. It is appropriate for small databases or any other environment that requires fault tolerance but small capacity.

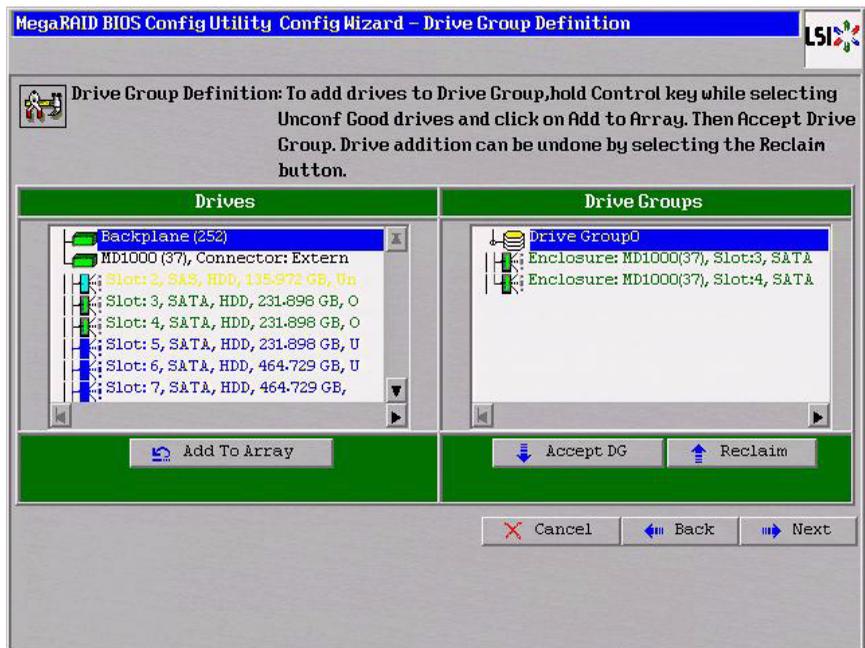
When you select **Manual Configuration** and click **Next**, the Drive Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while you select two ready drives in the Drives panel on the left. You must select an even number of drives.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in [Figure 3.11](#).

If you need to undo the changes, click the **Reclaim** button.

**Note:** A RAID 1 virtual drive can contain up to 16 drive groups and 32 drives in a single span. (Other factors, such as the type of controller, can limit the number of drives.) You must use two drives in each RAID 1 drive group in the span.

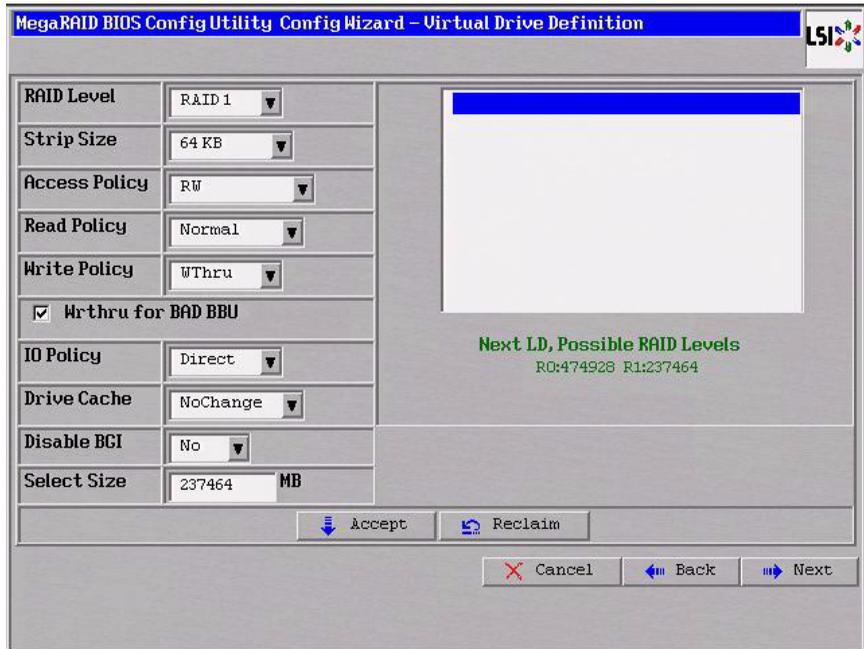
**Figure 3.11 WebBIOS Disk Group Definition Screen**



3. When you have finished selecting drives for the drive group, click **Accept DG**.
4. Click **Next**.

The Virtual Drive Definition screen appears, as shown in [Figure 3.12](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

**Figure 3.12 WebBIOS Virtual Drive Definition Screen**



5. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 1.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

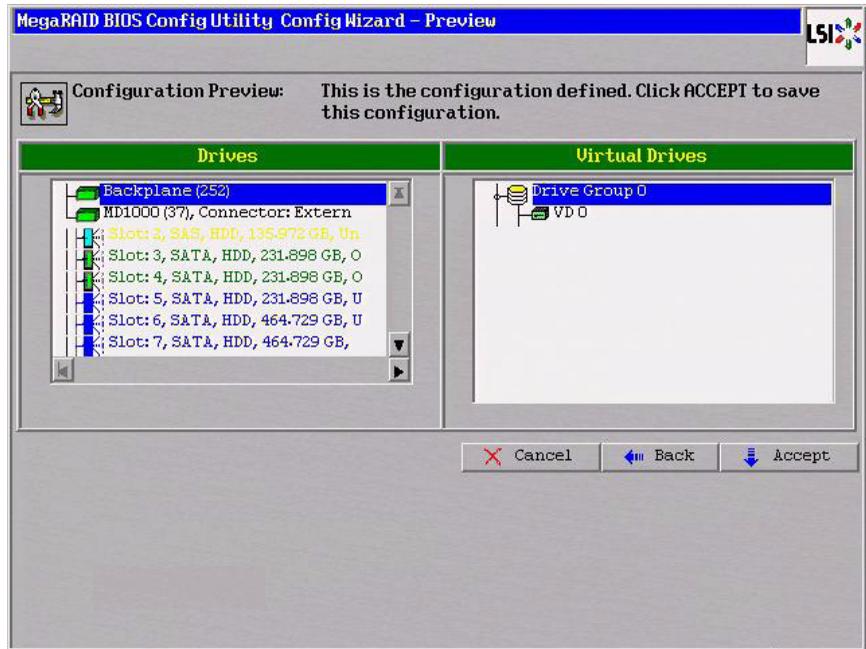
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access. This is the default.
  - ◇ *Read Only:* Allow read-only access.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Cache:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive(s) in megabytes. Normally, this would be the full size for RAID 1 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
6. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
  7. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.13](#).

**Figure 3.13 RAID 1 Configuration Preview**



8. Check the information in the configuration preview.
9. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
10. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

#### **3.4.3.4 Using Manual Configuration: RAID 5**

RAID 5 uses drive striping at the block level and parity. In RAID 5, the parity information is written to all drives. It is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously. RAID 5 provides data redundancy, high read rates, and good performance in most environments. It also provides redundancy with lowest loss of capacity.

RAID 5 provides high data throughput. RAID 5 is useful for transaction processing applications because each drive can read and write

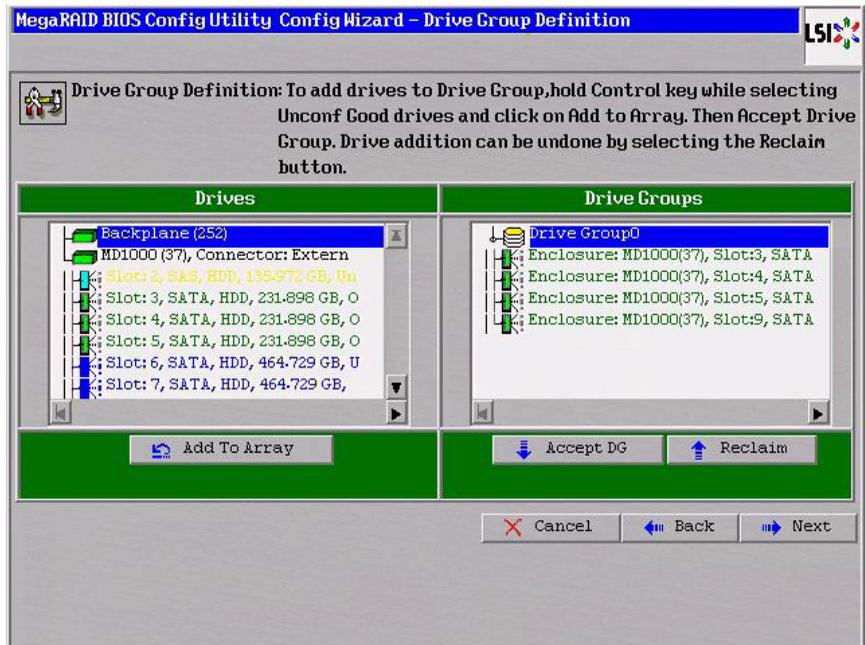
independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. You can use RAID 5 for office automation and online customer service that require fault tolerance. In addition, RAID 5 is good for any application that has high read request rates but low write request rates.

When you select **Manual Configuration** and click **Next**, the Drive Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while you select at least three ready drives in the Physical Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in [Figure 3.14](#).

If you need to undo the changes, click the **Reclaim** button.

**Figure 3.14 WebBIOS Disk Group Definition Screen**

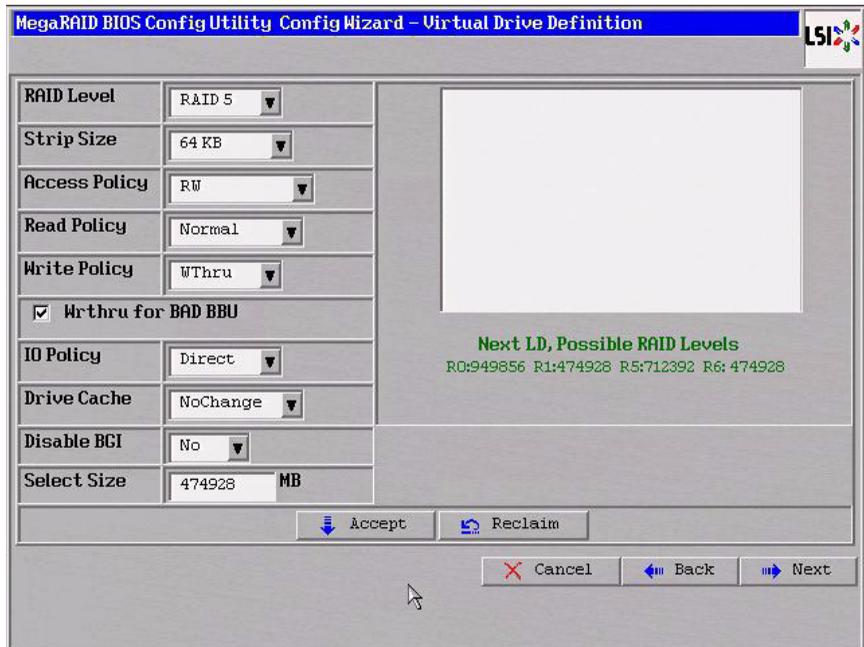


3. When you have finished selecting drives for the drive group, click **Accept DG**.

4. Click **Next**.

The Virtual Drive Definition screen appears, as shown in [Figure 3.15](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

**Figure 3.15 WebBIOS Virtual Drive Definition Screen**



5. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 5.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read

performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

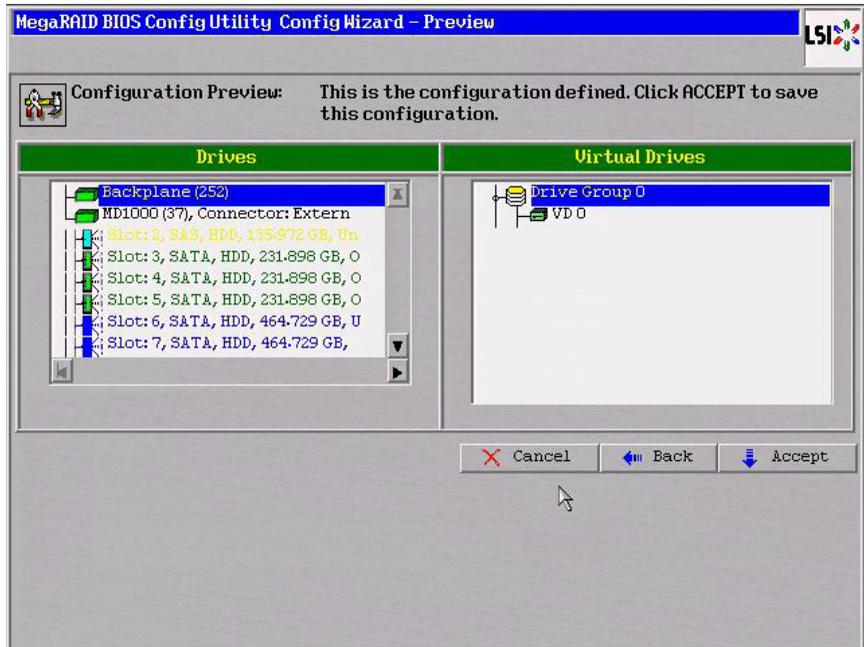
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access. This is the default.
  - ◇ *Read Only:* Allow read-only access.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Cache:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 5 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
6. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
  7. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.16](#).

**Figure 3.16 RAID 5 Configuration Preview**



8. Check the information in the configuration preview.
9. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation and return to the WebBIOS main menu, or click **Back** to return to the previous screens and change the configuration.
10. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

### 3.4.3.5 Using Manual Configuration: RAID 6

RAID 6 is similar to RAID 5 (drive striping and distributed parity), except that instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two drives in a virtual drive without losing data. Use RAID 6 for data that requires a very high level of protection from loss.

RAID 6 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously. It provides data redundancy, high read rates, and good performance in most environments.

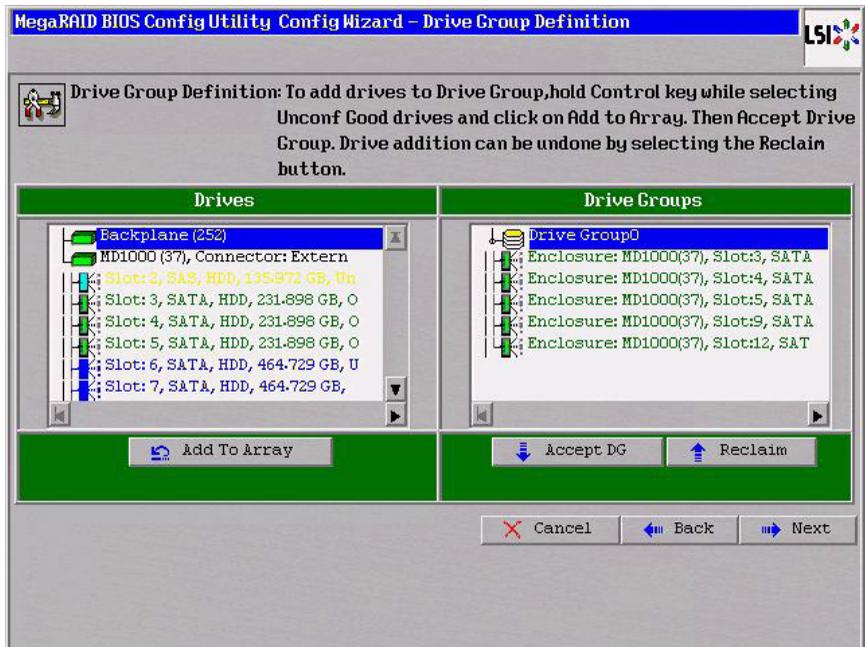
In the case of a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive, and then the other failed drive.

When you select **Manual Configuration** and click **Next**, the drive Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in [Figure 3.14](#).

If you need to undo the changes, click the **Reclaim** button.

Figure 3.17 WebBIOS Drive Group Definition Screen



3. When you have finished selecting drives for the drive group, click **Accept DG** for each.
4. Click **Next**.

The Virtual Drive Definition screen appears, as shown in [Figure 3.18](#). Use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

**Figure 3.18 WebBIOS Virtual Drive Definition Screen**

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level	RAID 6
Strip Size	64 KB
Access Policy	RW
Read Policy	Normal
Write Policy	WThru
<input checked="" type="checkbox"/> Wthru for BAD BBU	
IO Policy	Direct
Drive Cache	NoChange
Disable BGI	No
Select Size	712392 MB

Next LD, Possible RAID Levels  
R0:1187320 R5:949856 R6: 712392

Accept Reclaim

Cancel Back Next

5. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 6.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

**Note:** WebBIOS does not allow you to select 8 Kbytes as the stripe size when you create a RAID 6 drive group with three drives.

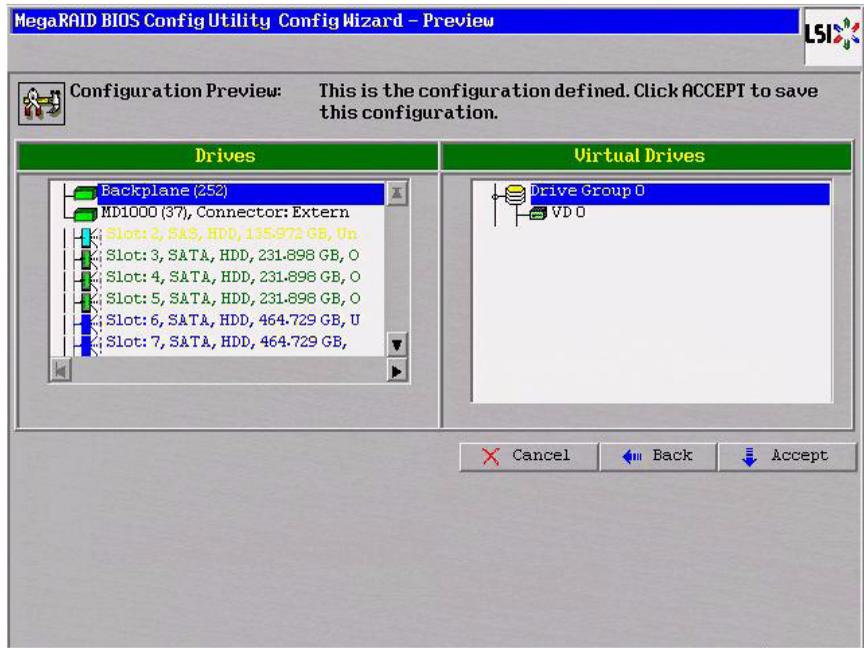
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access. This is the default.
  - ◇ *Read Only:* Allow read-only access.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Policy:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This drive policy is the default.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 6 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
6. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
  7. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.16](#).

**Figure 3.19 RAID 6 Configuration Preview**



8. Check the information in the configuration preview.
9. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
10. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

### 3.4.3.6 Using Manual Configuration: RAID 10

RAID 10, a combination of RAID 1 and RAID 0, has mirrored drives. It breaks up data into smaller blocks, then stripes the blocks of data to each RAID 1 drive group. Each RAID 1 drive group then duplicates its data to its other drive. The size of each block is determined by the stripe size parameter, which is 64 Kbytes. RAID 10 can sustain one drive failure in each drive group while maintaining data integrity.

RAID 10 provides both high data transfer rates and complete data redundancy. It works best for data storage that must have 100 percent

redundancy of RAID 1 (mirrored drive groups) and that also needs the enhanced I/O performance of RAID 0 (striped drive groups); it works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.

When you select **Manual Configuration** and click **Next**, the Drive Group Definition screen appears.

You use the Drive Group Definition screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting two ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed two-drive drive group configuration in the Drive Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

3. Click **Accept DG** to create a RAID 1 drive group.

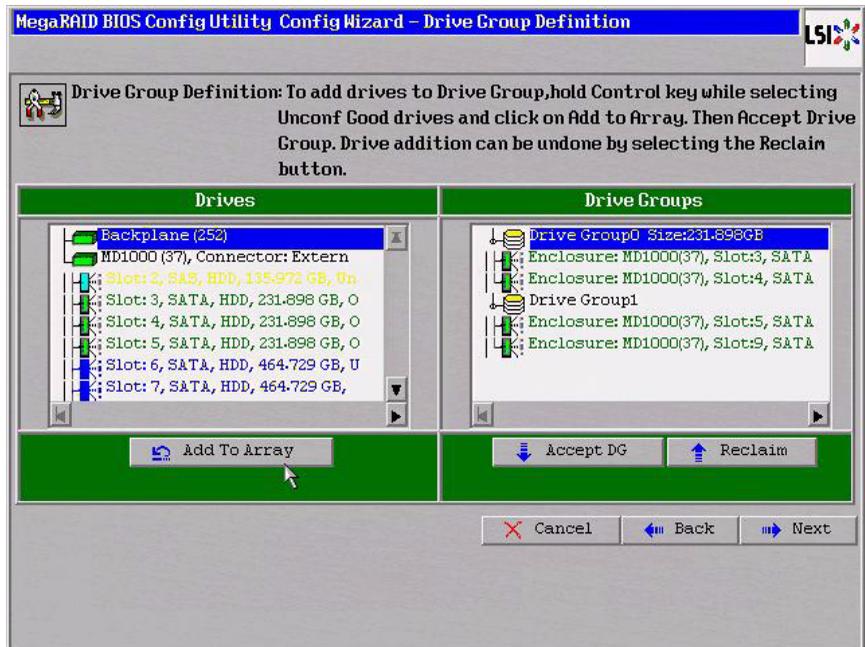
An icon for the next drive group displays in the right panel.

4. Click on the icon for the next drive group to select it.
5. Hold <Ctrl> while selecting two more ready drives in the Drives panel to create a second RAID 1 drive group with two drives.
6. Click **Add To Array** to move the drives to a second two-drive drive group configuration in the Drive Groups panel, as shown in [Figure 3.20](#).

If you need to undo the changes, click the **Reclaim** button.

**Note:** RAID 10 supports a maximum of eight spans, with a maximum of 32 drives per span. (Other factors, such as the type of controller, can limit the number of drives.) You must use an even number of drives in each RAID 10 drive group in the span.

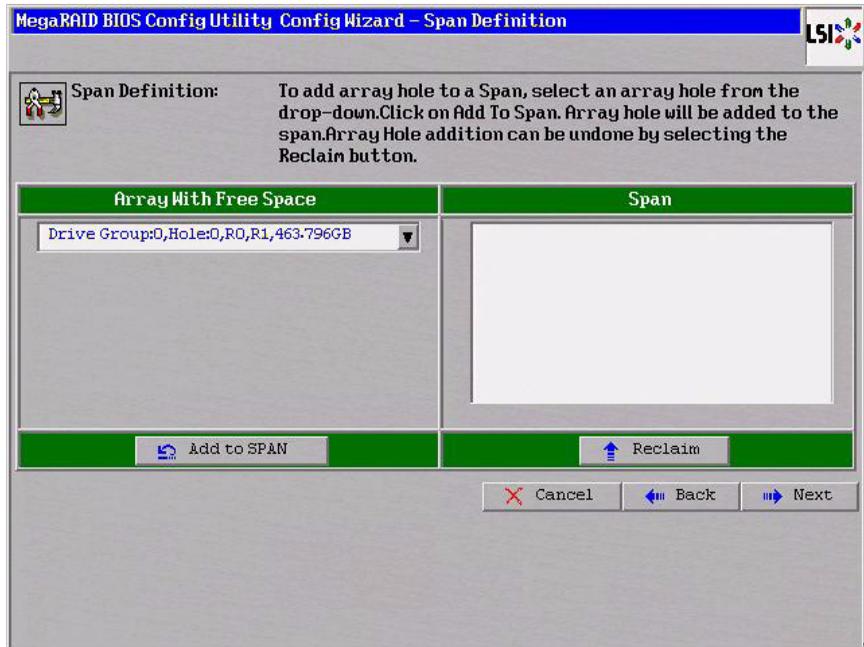
Figure 3.20 WebBIOS Drive Group Definition Screen



7. Repeat [step 4](#) to [step 6](#) until you have selected all the drives you want for the drive groups.
8. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
9. Click **Next**.

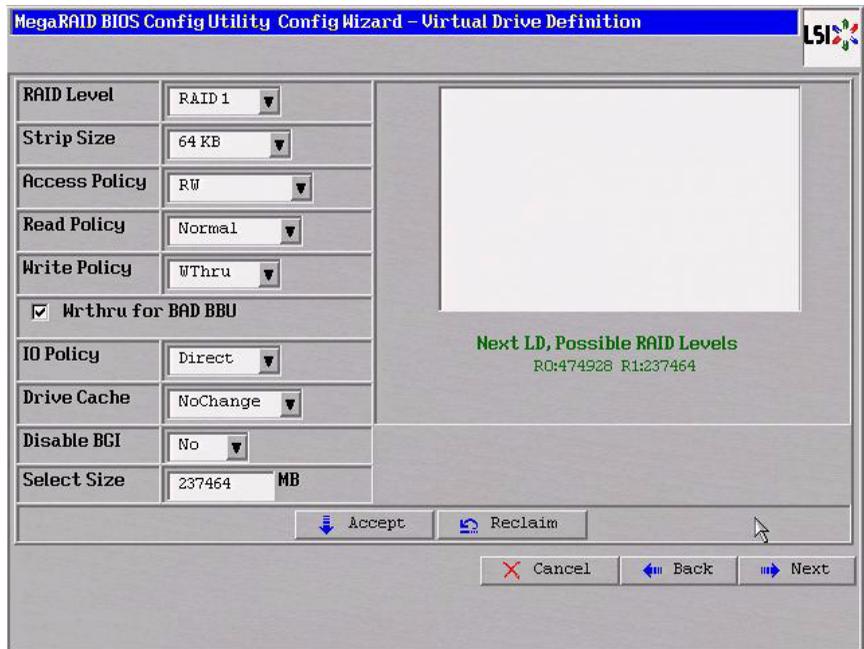
The Span Definition screen appears, as shown in [Figure 3.21](#). This screen displays the drive group holes you can select to add to a span.

**Figure 3.21 WebBIOS Span Definition Screen**



10. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group with two drives, and click **Add to SPAN**.  
The drive group you select displays in the right frame under the heading **Span**.
11. Hold <Ctrl> while you select a second drive group with two drives, and click **Add to SPAN**.  
Both drive groups display in the right frame under **Span**.
12. If there are additional drive groups with two drives each, you can add them to the virtual drive.
13. Click **Next**.  
The Virtual Drive Definition screen appears, as shown in [Figure 3.22](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.
14. Hold <Ctrl> while you select two drive groups with two drives in the Configuration panel on the right.

**Figure 3.22 WebBIOS Virtual Drive Definition Screen**



**Note:** The WebBIOS Configuration Utility displays the maximum available capacity while creating the RAID 10 drive group. In version 1.03 of the utility, the maximum size of the RAID 10 drive group is the sum total of the two RAID 1 drive groups. In version 1.1, the maximum size is the size of the smaller drive group multiplied by two.

15. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 10.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read

performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

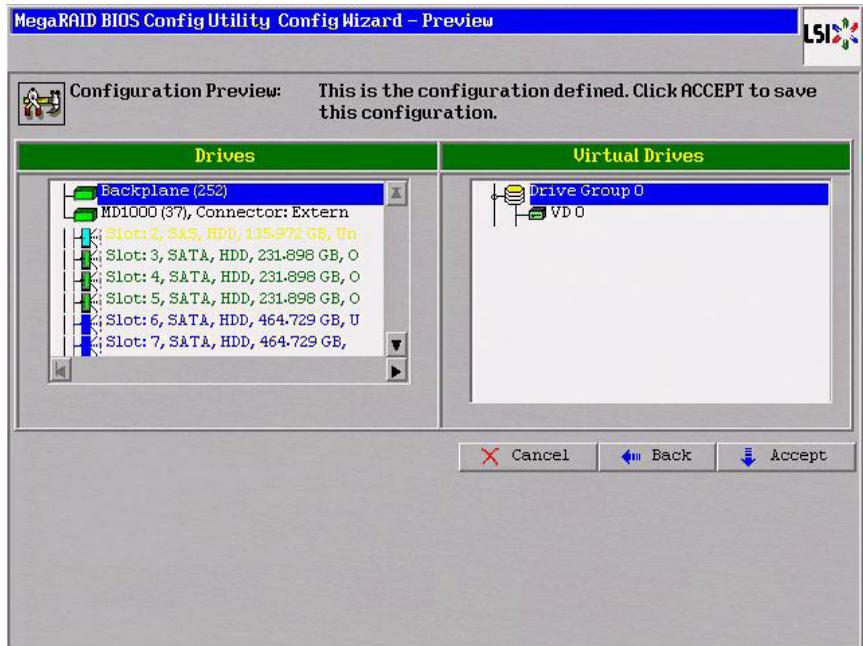
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access.
  - ◇ *Read Only:* Allow read-only access. This is the default.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Policy:** Specify the drive cache policy:
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 10 shown in the configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
16. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
17. When you are finished defining virtual drives, click **Next** .

The Configuration Preview screen appears, as shown in [Figure 3.23](#).

**Figure 3.23 RAID 10 Configuration Preview**



18. Check the information in the configuration preview.
19. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation and return to the WebBIOS main menu, or click **Back** to return to the previous screens and change the configuration.
20. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

### 3.4.3.7 Using Manual Configuration: RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 uses both distributed parity and drive striping across multiple drive groups. It provides high data throughput, data redundancy, and very good performance. It is best implemented on two RAID 5 drive groups with data striped across both drive groups. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

RAID 50 is appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.

When you select **Manual Configuration** and click **Next**, the Drive Group Definition screen appears. You use this screen to select drives to create drive group.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

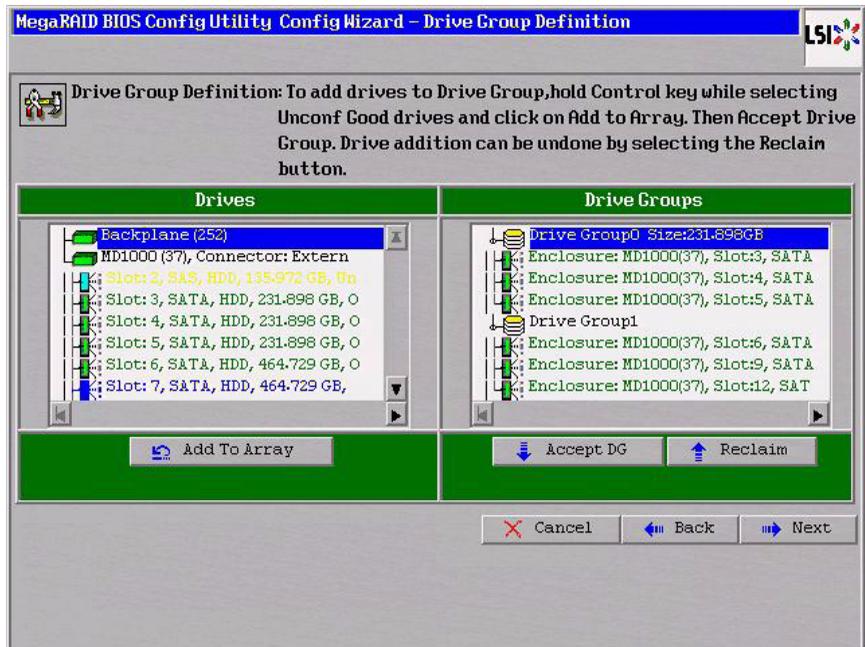
3. Click **Accept DG** to create a RAID 5 drive group.

An icon for a second drive group displays in the right panel.

4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in [Figure 3.24](#).

If you need to undo the changes, click the **Reclaim** button.

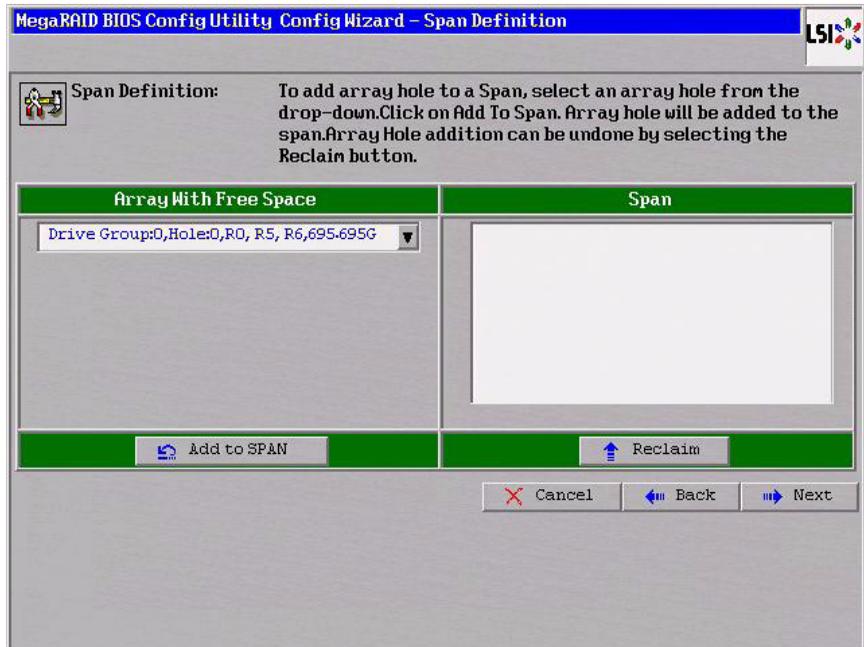
Figure 3.24 WebBIOS Disk Group Definition Screen



7. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
8. Click **Next**.

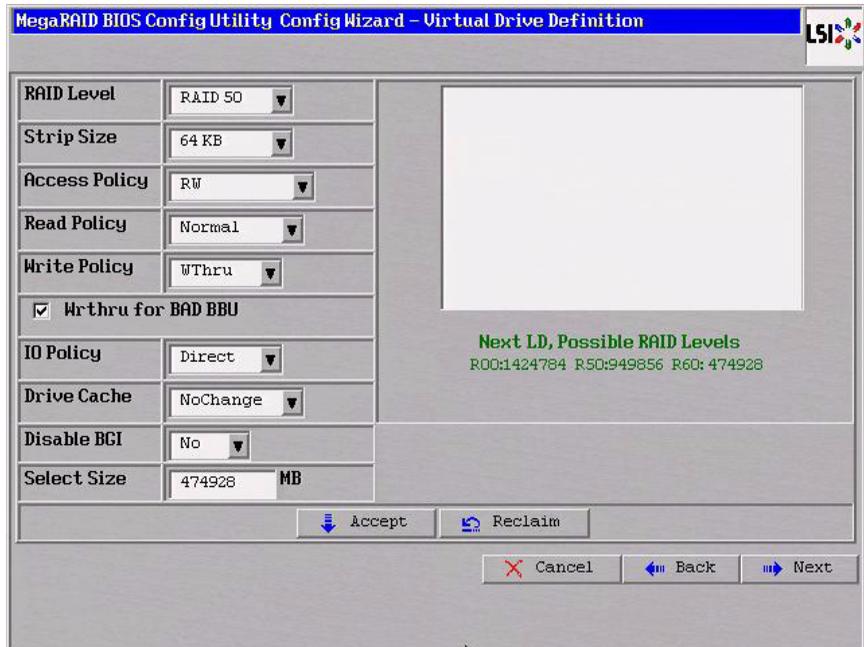
The Span Definition screen appears, as shown in [Figure 3.25](#). This screen displays the drive group holes you can select to add to a span.

**Figure 3.25 WebBIOS Span Definition Screen**



9. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group of three or more drives, and click **Add to SPAN**.  
The drive group you select displays in the right frame under the heading **Span**.
10. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.  
Both drive groups display in the right frame under **Span**.
11. Click **Next**.  
The Virtual Drive Definition screen appears, as shown in [Figure 3.26](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drive(s).
12. Hold <Ctrl> while you select two 3-drive drive groups in the Configuration panel on the right.

**Figure 3.26 WebBIOS Virtual Drive Definition Screen**



13. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 50.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

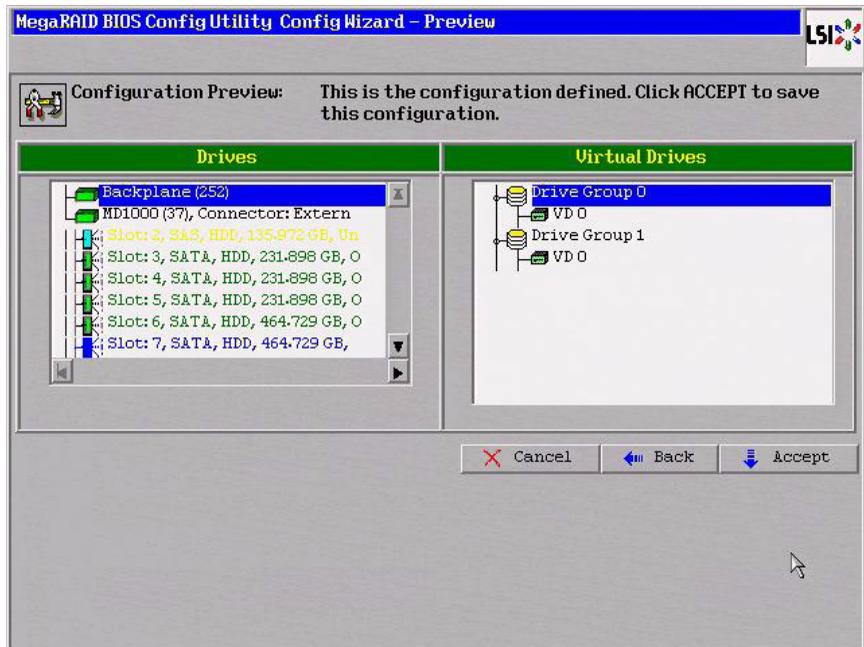
- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access.
  - ◇ *Read Only:* Allow read-only access. This is the default.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
    - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
    - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
  - **Drive Policy: Specify the drive cache policy:**
    - ◇ *Enable:* Enable the drive cache.
    - ◇ *Disable:* Disable the drive cache. This drive policy is the default.
    - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
  - **Disable BGI:** Specify the background initialization status:
    - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
    - ◇ *Yes:* Select **Yes** if you do not want to allow background initializations for configurations on this controller.
  - **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 50 shown in the Configuration Panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
14. Click **Accept** to accept the changes to the virtual drive definition or click **Reclaim** to return to the previous settings.
15. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.27](#).

Figure 3.27 RAID 50 Configuration Preview



16. Check the information in the configuration preview.
17. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
18. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

#### 3.4.3.8 Using Manual Configuration: RAID 60

RAID 60 provides the features of both RAID 0 and RAID 6, and includes both parity and drive striping across multiple drive groups. RAID 6 supports two independent parity blocks per stripe. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. RAID 60 is best implemented on two RAID 6 drive groups with data striped across both drive groups. Use RAID 60 for data that requires a very high level of protection from loss.

RAID 60 can support up to eight spans and tolerate up to 16 drive failures, though less than total drive capacity is available. Two drive failures can be tolerated in each RAID 6 level drive group.

RAID 60 is appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

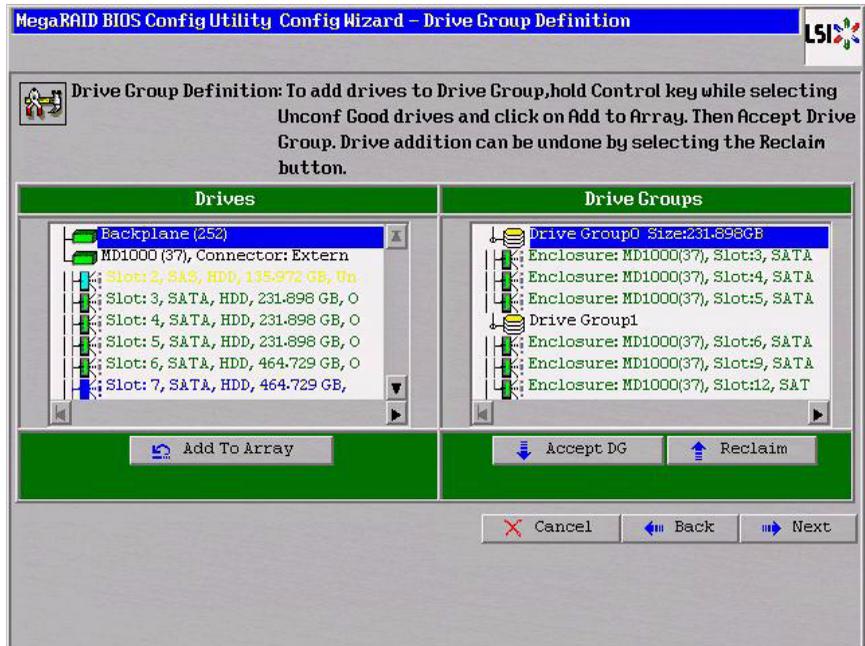
3. Click **Accept DG** to create a RAID 6 drive group.

An icon for a second drive group displays in the right panel.

4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in [Figure 3.24](#).

If you need to undo the changes, click the **Reclaim** button.

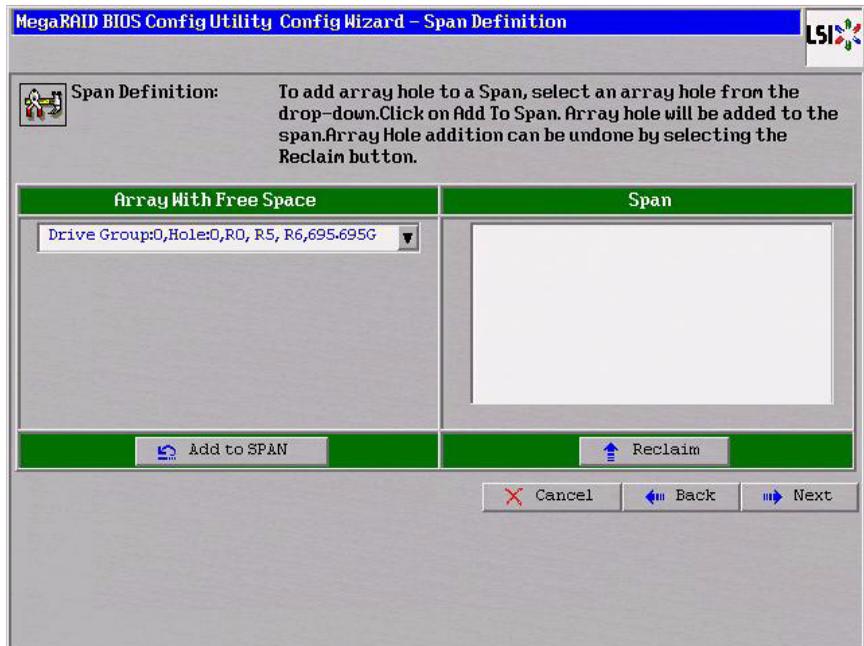
Figure 3.28 WebBIOS Drive Group Definition Screen



7. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
8. Click **Next**.

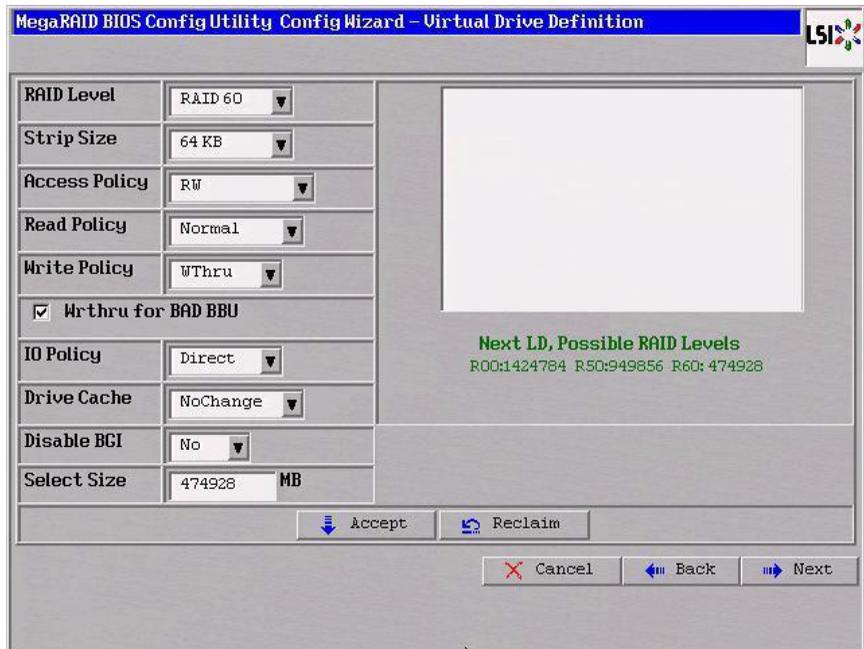
The Span Definition screen appears, as shown in [Figure 3.29](#). This screen displays the drive group holes you can select to add to a span.

**Figure 3.29 WebBIOS Span Definition Screen**



9. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group of three or more drives, and click **Add to SPAN**.  
The drive group you select displays in the right frame under the heading **Span**.
10. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.  
Both drive groups display in the right frame under **Span**.
11. Click **Next**.  
The Virtual Drive Definition screen appears, as shown in [Figure 3.26](#). You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drive(s).
12. Hold <Ctrl> while you select two 3-drive drive groups in the Configuration window on the right.

**Figure 3.30 WebBIOS Virtual Drive Definition Screen**



13. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 60.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
  - ◇ *RW:* Allow read/write access.
  - ◇ *Read Only:* Allow read-only access. This is the default.
  - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
  - ◇ *Normal:* This disables the read ahead capability. This is the default.
  - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
  - ◇ *Adaptive:* When Adaptive read ahead is selected, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to *Normal* (no read ahead).
- **Write Policy:** Specify the write policy for this virtual drive:
  - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
  - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
  - ◇ *Bad BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution:** LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use **either** a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

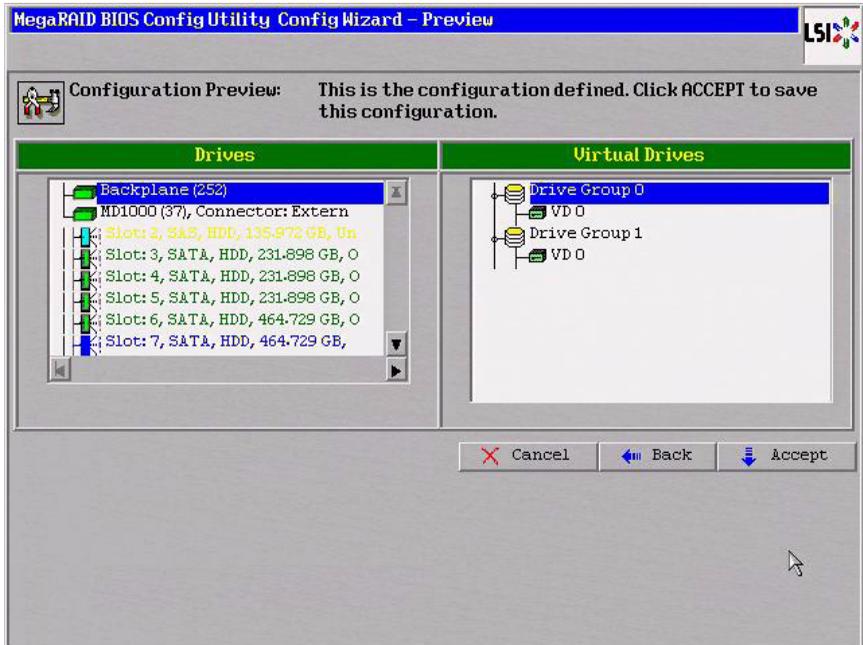
- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
  - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
  - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
- **Drive Cache: Specify the drive cache policy:**
  - ◇ *Enable:* Enable the drive cache.
  - ◇ *Disable:* Disable the drive cache. This drive policy is the default.
  - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default drive cache policy.
- **Disable BGI:** Specify the background initialization status:
  - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
  - ◇ *Yes:* Select *Yes* if you do not want to allow background initializations for configurations on this controller.
- **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 60 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.

**Note:** WebBIOS does not allow you to select 8 Kbytes as the stripe size when you create a RAID 60 drive group with six drives.

- Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
- Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in [Figure 3.27](#).

**Figure 3.31 RAID 60 Configuration Preview**



- Check the information in the configuration preview.
- If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, or click **Back** to return to the previous screens and change the configuration.
- If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

## 3.5 Viewing and Changing Device Properties

This section explains how you can use the WebBIOS CU to view and change the properties for controllers, virtual drives, drives, and BBUs.

### 3.5.1 Viewing and Changing Controller Properties

WebBIOS displays information for one LSI SAS controller at a time. If your computer system has multiple LSI SAS controllers, you can view information for a different controller by clicking **Controller Selection** on the main screen. When the Controller Selection screen appears, select the controller you want from the list.

To view the properties for the currently selected controller, click **Controller Properties** on the main WebBIOS screen. There are three Controller Properties screens. [Figure 3.32](#) shows the first screen.

**Figure 3.32 First Controller Properties Screen**



MegaRAID SAS 8888ELP			
Serial Number	123456	FRU	None
SubVendorID	0x1000	Encryption Capable	No
SubDeviceID	0x1006	NURAMSize	32 KB
PortCount	8	Memory Size	256 MB
HostInterface	PCI-E	Min Stripe Size	8 KB
Firmware Version	1.40-02-0514	Max Stripe Size	1024 KB
FW Package Version	9.1.1-0012	Virtual Drive Count	0
Firmware Time	Aug 29 2006:18:40:33	Drive Count	14
WebBIOS Version	2.2-13-Rel		

Next

Home Back

The information on this screen is read-only and cannot be modified directly. Most of this information is self-explanatory. The screen lists the number of virtual drives that are already defined on this controller, and the number of drives connected to the controller.

If a background initialization is in progress, you can click **Background Init Progress** to determine its state of completion. Click **Next** to view the second Controller Properties screen, as shown in [Figure 3.33](#).

Figure 3.33 Second Controller Properties Screen

Properties			
Battery Backup	None	Coercion Mode	1GB-way
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Enabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	2
Reconstruction Rate	30	Spinup Delay	12
Controller BIOS	Enabled	StopOnError	Enabled
NCQ	Enabled	Drive Powersave	Disabled
Connector 1	External	Connector 2	External

Submit    Reset    Next

Home    Back

Click **Next** to view the third Controller Properties screen, as shown in Figure 3.33.

**Figure 3.34 Third Controller Properties Screen**

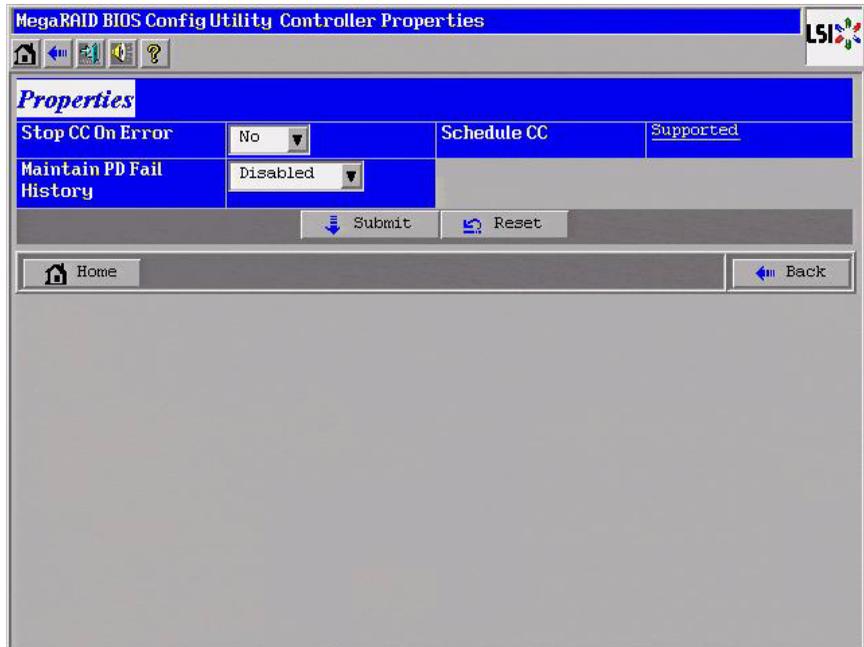


Table 3.2 describes the entries/options listed on the second and third Controller Properties screen. LSI recommends that you leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.

**Table 3.2 Controller Properties Menu Options**

Option	Description
Battery Backup	This entry indicates whether the selected controller has a BBU. If present, you can click <i>Present</i> to view information about the BBU. For more information, see <a href="#">Section 3.5.4, “Viewing Battery Backup Unit Information.”</a>
Set Factory Defaults	Use this option to load the default MegaRAID® WebBIOS CU settings. The default is <i>No</i> .
Cluster Mode	Use this option to enable or disable Cluster mode. The default is <i>Disabled</i> . A cluster is a grouping of independent servers that can access the same data storage and provide services to a common set of clients. When Cluster mode is disabled, the system operates in Standard mode.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The default is 30 percent. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources devoted to a rebuild.

**Table 3.2 Controller Properties Menu Options (Cont.)**

Option	Description
BGI Rate	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The default is 30 percent.
CC Rate	Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The default is 30 percent.
Reconstruction Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The default is 30 percent.
Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. The default is <i>Enabled</i> . If the boot device is on the selected controller, the BIOS must be enabled; otherwise, the BIOS should be disabled or it might not be possible to use a boot device elsewhere.
NCQ	Native Command Queuing (NCQ) gives an individual drive the ability to optimize the order in which it executes the read and write commands. The default is <i>Enabled</i> .
Connector 1	Identifies where the chain of enclosures is connected to the RAID controller.
Coercion Mode	Drive coercion is a tool for forcing drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are <i>None</i> , <i>128MB-way</i> , and <i>1GB-way</i> . The default is <i>None</i> .  Note: The number you choose depends on how much the drives from various vendors vary in their actual size. LSI recommends that you use the 1GB coercion mode option.
S.M.A.R.T. Polling	Use this option to determine how frequently the controller polls for drives reporting a Predictive Drive Failure (S.M.A.R.T.: Self-Monitoring Analysis and Reporting Technology error). The default is 300 seconds (5 minutes).
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller. The default is <i>Disabled</i> .
Patrol Read Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The default is 30 percent. The patrol read rate is the percentage of system resources dedicated to running a patrol read. See <a href="#">Section 4.5, "Patrol Read-Related Controller Properties"</a> for additional information about patrol read.
Cache Flush Interval	Use this option to control the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
Spinup Drive Count	Use this option to control the number of drives that spin up simultaneously. The default is 2 drives.
Spinup Delay	Use this option to control the interval (in seconds) between spinup of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The default is 12 seconds.

**Table 3.2 Controller Properties Menu Options (Cont.)**

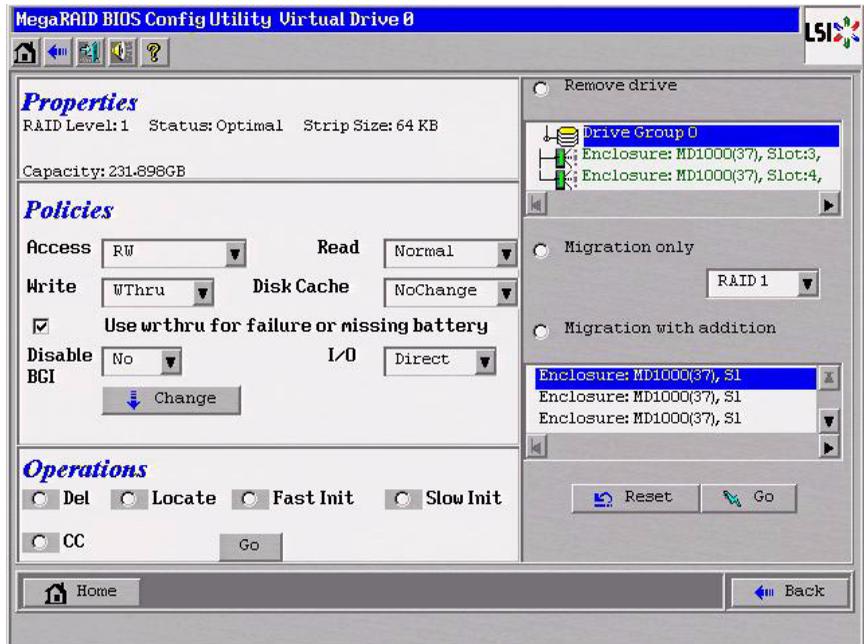
<b>Option</b>	<b>Description</b>
StopOnError	Enable this option if you want the boot process to stop when the controller BIOS encounters an error during boot-up. The default is <i>Disabled</i> .
Drive Powersave	Drive Powersave conserves energy by placing certain unused drives into powersave mode. Use this field to choose whether to allow unconfigured drives to enter powersave mode.  When this option is selected, unconfigured drives may be spun down. When not selected, these drives are not spun down. The controller will automatically spin up drives from powersave mode whenever necessary. The powersave option is not selected by default. You have to select it to enable the spin-down of drives.
Connector 2	Identifies where the chain of enclosures is connected to the RAID controller.
Stop CC on Error	Enable this option if you want to stop a consistency check when the controller BIOS encounters an error. The default is <i>No</i> .
Maintain PD Fail History	Enable this option to maintain the history of all drive failures. The default is <i>Disabled</i> .
Schedule CC	Indicates whether the option to schedule the date and time for a consistency check is supported.

If you make changes to the options on this screen, click **Submit** to register them. If you change your mind, click **Reset** to return the options to their default values.

### **3.5.2 Viewing and Changing Virtual Drive Properties**

Access the Virtual Drive screen by clicking on a virtual drive in the list of virtual drives in the right panel on the WebBIOS CU main screen. The Virtual Drive screen displays, as shown in [Figure 3.35](#).

Figure 3.35 Virtual Drive Screen



The Properties panel of this screen displays the virtual drive's RAID level, state, size, and stripe size.

The Policies panel lists the virtual drive policies that were defined when the storage configuration was created. For information about these policies, see [Section 3.4.3, “Using Manual Configuration.”](#) To change any of these policies, make a selection from the drop-down menu and click **Change**.

The Operations panel lists operations that can be performed on the virtual drive. To perform an operation, select it and click **Go**. Then choose from the following options:

- Select **Del** to delete this virtual drive. For more information, see [Section 3.7.2, “Deleting a Virtual Drive.”](#)
- Select **Locate** to make the LEDs flash on the drives used by this virtual drive. This works only if the drives are installed in a drive enclosure that supports SAFTE.
- Select **Fast Init** or **Slow Init** to initialize this virtual drive. A fast initialization quickly writes zeroes to the first and last 10 Mbyte

regions of the new virtual drive and then completes the initialization in the background. A slow initialization is not complete until the entire virtual drive has been initialized with zeroes. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.

**Caution:** Before you run an initialization, back up any data on the virtual drive that you want to save. All data on the virtual drive is lost when you initialize it.

- Select **CC** to run a consistency check on this virtual drive. For more information, see [Section 3.7.1, “Running a Consistency Check.”](#) (This option is not available for RAID 0 virtual drives.)

In the right panel of the Virtual Drive screen you can change the virtual drive configuration by adding or removing a drive or by changing the RAID level.

**Caution:** Before you change a virtual drive configuration, back up any data on the virtual drive that you want to save.

To remove a drive from a virtual drive, select the drive in the small panel beneath the *Remove drive* option. Then select **Remove drive** and click **Go** at the bottom of the panel.

See [Section 3.7.4, “Migrating the RAID Level of a Virtual Drive”](#) for information about adding a drive to a virtual drive or migrating its RAID level.

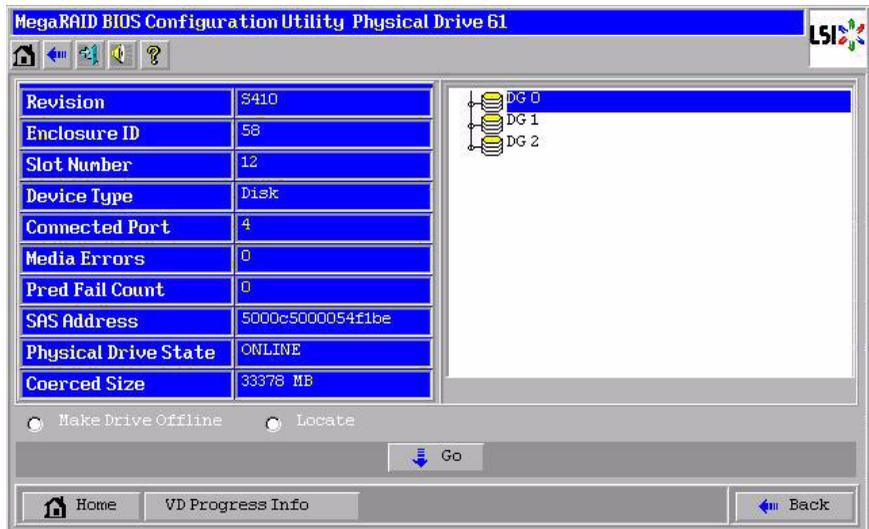
### 3.5.3 Viewing Drive Properties

The Physical Drive screen displays the properties of a selected drive and enables you to perform operations on the drive. There are two ways to access the Physical Drive screen:

- On the main menu screen, click on a drive in the right panel under the heading **Physical Drives**.
- On the main menu screen, click on **Physical Drives** in the left panel to display the Physical Drive screen. Then click on a drive in the right panel. Click on the **Properties** button, and click **Go**. The properties for the selected drive displays.

[Figure 3.36](#) shows the Physical Drive screen.

**Figure 3.36 Physical Drive Screen**



The drive properties are view-only and are self-explanatory. Note that the properties include the state of the drive.

Operations you can perform are listed at the bottom of the screen. After you select an operation, click **Go** to start the operation. The operations vary depending on the drive state. If the drive state is **Online**, the following operations appear:

- Select **MakeDriveOffline** if you want to force the drive offline.

Note: If you force offline a good drive that is part of a redundant drive group with a hot spare, the drive will rebuild to the hot spare drive. The drive you forced offline will go into the *Unconfigured Bad* state. Access the BIOS utility to set the drive to the *Unconfigured Good* state.

- Select **Locate** to make the LED flash on the drive. This works only if the drive is installed in a drive enclosure.

If the drive state is Unconfigured Good, four additional operations appear on this screen:

- Select **Make Global HSP** to make a global hot spare, available to all of the virtual drives.

- Select **Make Dedicated HSP** to make a hot spare dedicated to a specific virtual drive.

WebBIOS displays the global hot spare as `Global` and the dedicated hot spare as `Ded`. The icon for the dedicated hot spare displays under its associated virtual drive. The drive number, drive state, drive capacity, and drive manufacturer display.

- Select **Enclosure Affinity** so if there are drive failures present on a split backplane configuration, then the hot spare will be used first on the backplane side that it resides in.
- Select **Prepare for Removal** to prepare the drive for removal from the enclosure.

The **Prepare for Removal** feature is different from spinning a drive down into powersave mode because it also involves flagging the drive as ready to remove. Therefore, if you choose to prepare a drive for removal, **Ready to Remove** displays in the device tree for that drive, instead of **Powersave**.

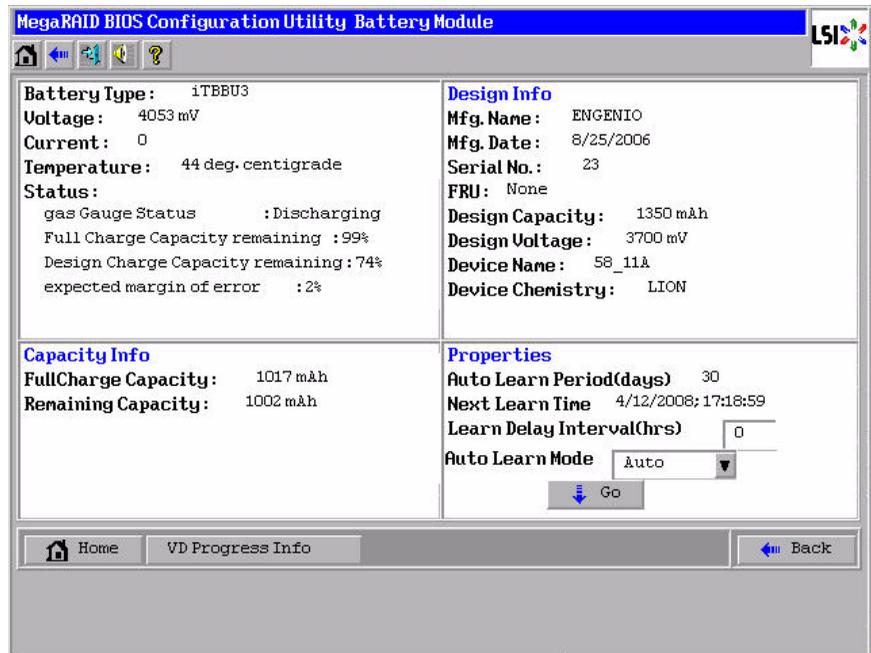
### 3.5.4 Viewing Battery Backup Unit Information

If your SAS controller has a battery backup unit (BBU), you can view information about it. To do this, follow these steps:

1. Click **Adapter Properties** on the main WebBIOS CU screen.
2. Click **Next** to view the second Adapter Properties screen.
3. Click the word *Present* in the Battery Backup field at the top left of the screen.

The Battery Module screen appears, as shown in [Figure 3.37](#).

Figure 3.37 Battery Module Screen



Most of the Battery Module properties are view-only and are self-explanatory. In the lower right panel, there are two properties that can be changed. (LSI recommends that you leave these properties at their default settings.)

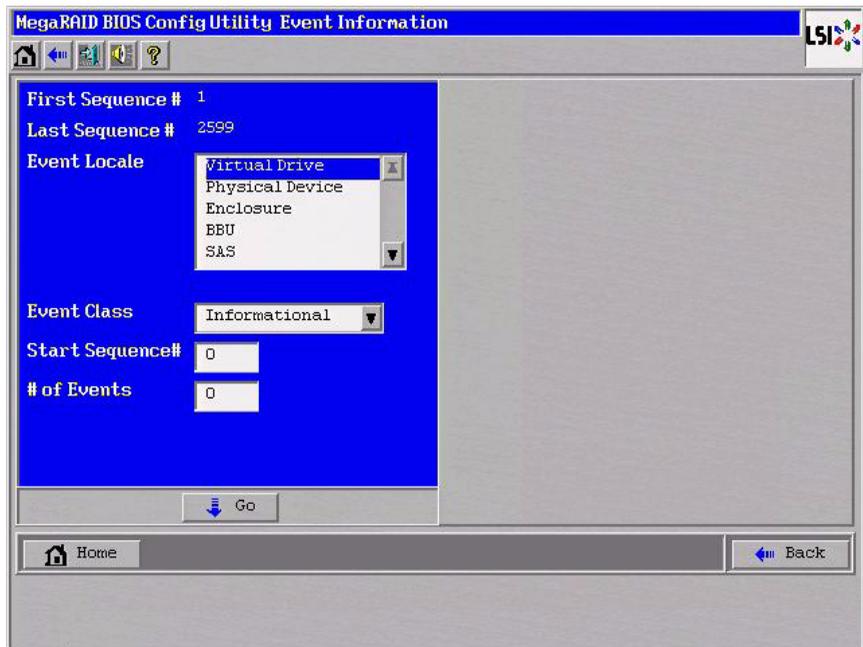
A *learning cycle* is a battery calibration operation performed by the controller periodically to determine the condition of the battery. To change the length of the interval between learning cycles, enter a different number of hours for Learn Delay Interval and click **Go**.

**Note:** LSI recommends 30 days for the interval between cycles.

## 3.6 Viewing System Event Information

The SAS controller firmware monitors the activity and performance of all storage configurations and devices in the system. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message is generated and is stored in the controller NVRAM. You can use the WebBIOS CU to view these event messages. To do this, click **Events** on the main WebBIOS CU screen. The Event Information screen appears, as shown in [Figure 3.38](#).

**Figure 3.38** Event Information Screen



The right side of the screen is blank until you select an event to view. The First Sequence and Last Sequence fields in the upper left of the screen show you how many event entries are currently stored.

To view event information, follow these steps:

1. Select an Event Locale from the menu. For example, select **Enclosure** to view events relating to the drive enclosure.
2. Select an Event Class: *Information*, *Warning*, *Critical*, *Fatal*, or *Dead*.

3. Enter a Start Sequence number, between the First Sequence and Last Sequence numbers. The higher the number, the more recent the event.
4. Enter the Number of events of this type that you want to view, and click **Go**.  
The first event in the sequence appears in the right panel.
5. Click **Next** or **Prev** to page forward or backward through the sequence of events.
6. If you want, select different event criteria in the left panel, and click **Go** again to view a different sequence of events.

Each event entry includes a timestamp and a description to help you determine when the event occurred and what it was.

---

## 3.7 Managing Configurations

This section includes information about maintaining and managing storage configurations.

### 3.7.1 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives. A consistency check verifies that the redundancy data is correct and available for RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 drive groups. To do this, follow these steps:

1. On the main WebBIOS CU screen, select a virtual drive.
2. Click **Virtual Drives**.
3. When the Virtual Drive screen appears, select **CC** in the lower left panel, and click **Go**.

The consistency check begins.

If the WebBIOS CU finds a difference between the data and the parity value on the redundant drive group, it assumes that the data is accurate and automatically corrects the parity value. Be sure to back up the data before running a consistency check if you think the consistency data may be corrupted.

## 3.7.2 Deleting a Virtual Drive

You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The WebBIOS CU provides a list of configurable drive groups where there is a space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the whole drive group.

To delete a virtual drive, follow these steps:

**Caution:** Back up any data that you want to keep before you delete the virtual drive.

1. On the main WebBIOS CU screen, select a virtual drive.
2. Click **Virtual Drives**.
3. When the Virtual Drive screen appears, select **Del** in the lower left panel, and click **Go**.
4. When the message appears, confirm that you want to delete the virtual drive.

## 3.7.3 Importing or Clearing a Foreign Configuration

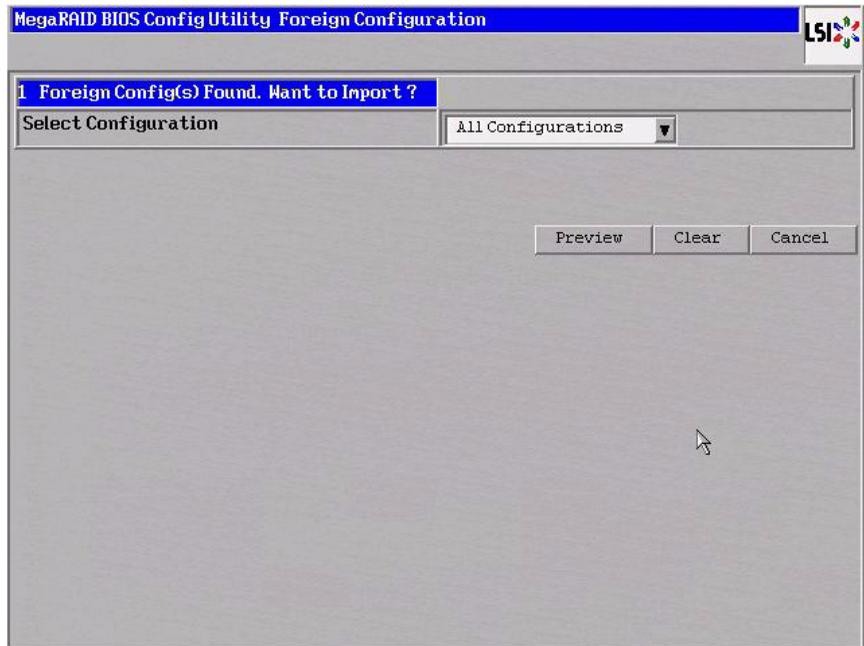
A *foreign configuration* is a storage configuration that already exists on a replacement set of drives that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

The BIOS CU allows you to import the foreign configuration to the RAID controller, or to clear the configuration so you can create a new configuration using these drives.

**Note:** When you create a new configuration, the WebBIOS CU shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, will **not** appear. To use drives with existing configurations, you must first clear the configuration on those drives.

If WebBIOS CU detects a foreign configuration, the import screen appears, as shown in [Figure 3.39](#).

**Figure 3.39 Foreign Configuration Import Screen**



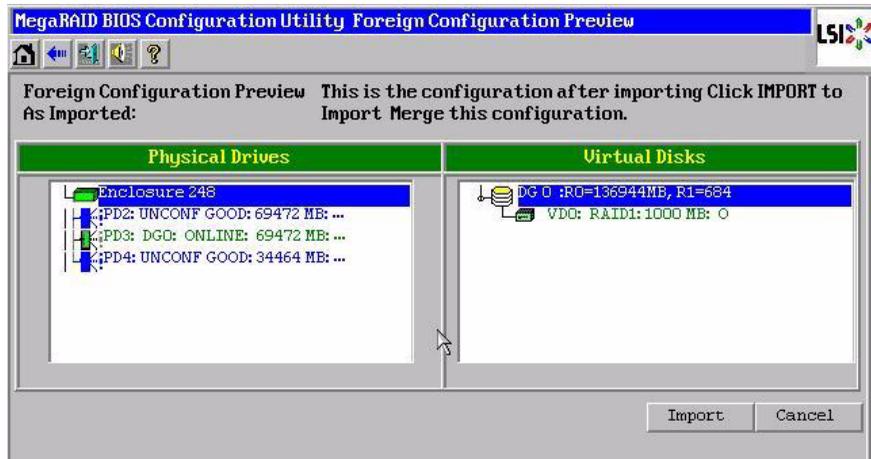
The GUID (Global Unique Identifier) entries on the drop-down list are OEM names and will vary from one installation to another.

Click **Preview** if you want to preview the foreign configuration. The preview screen appears, as shown in [Figure 3.40](#).

Click **Clear** if you want to clear the configuration and reuse the drives for another virtual drive.

Click **Cancel** to cancel the importation or preview of the configuration.

**Figure 3.40 Foreign Configuration Preview Screen**



The right panel shows the virtual drive properties of the foreign configuration. In this example, there is a RAID 1 virtual drive with 1,000 Mbytes. The left panel shows the drives that comprise the foreign configuration.

Click **Import** to import this foreign configuration and use it on this controller.

Click **Cancel** to clear the configuration and reuse the drives for another virtual drive.

### 3.7.3.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

If one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

Use the **Foreign Configuration Preview** screen to import or clear the foreign configuration in each case. The import procedure and clear procedure are described in [Section 3.7.3, “Importing or Clearing a Foreign Configuration.”](#)

The following scenarios can occur with cable pulls or drive removals.

**Note:** If you want to import the foreign configuration in any of the following scenarios, you should have all of the drives in the enclosure before you perform the import operation.

1. Scenario #1: If all of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds will occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 3.7.1, "Running a Consistency Check,"](#) for more information about checking data consistency.

2. Scenario #2: If some of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds will occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 3.7.1, "Running a Consistency Check,"](#) for more information about checking data consistency.

3. Scenario #3: If all of the drives in a virtual drive are removed, but at different times, and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, all drives that were pulled *before* the virtual drive became offline will be imported and then automatically rebuilt. Automatic rebuilds will occur in redundant virtual drives.

4. If the drives in a non-redundant virtual drive are removed, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. No rebuilds will occur after the import operation because there is no redundant data to rebuild the drives with.

### 3.7.3.2 Importing Foreign Configurations from Integrated RAID to MegaRAID

The LSI Integrated RAID solution simplifies the configuration options and provides firmware support in its host controllers. LSI offers two types of Integrated RAID (IR): Integrated Mirroring (IM) and Integrated Striping (IS).

You can import an IM or IS RAID configuration from an IR system into a MegaRAID system. The MegaRAID system treats the IR configuration as a foreign configuration. You can import or clear the IR configuration.

**Note:** For more information about Integrated RAID, refer to the *Integrated RAID for SAS User's Guide*. You can find this document on the LSI web site at:  
<http://www.lsi.com/cm/DownloadSearch.do>.

### 3.7.3.3 Troubleshooting Information

An IR virtual drive can have either 64 Mbytes or 512 Mbytes available for metadata at the end of the drive. This data is in LSI Data Format (LDF). MegaRAID virtual drives have 512 Mbytes for metadata at the end of the drive in the Disk Data format (DDF).

To import an IR virtual drive into MegaRAID, the IR virtual drive must have 512 Mbytes in the metadata, which is the same amount of metadata as in a MegaRAID virtual drive. If the IR virtual drive has only 64 Mbytes when you attempt to import it into MegaRAID, the import will fail because the last 448 Mbytes of your data will be overwritten and the data lost.

If your IR virtual drive has only 64 Mbytes for metadata at the end of the drive, you cannot import the virtual drive into MegaRAID. You need to use another upgrade method, such as backup/restore to the upgraded virtual drive type.

In order to import an IR virtual drive into a MegaRAID system, use the **Foreign Configuration Preview** screen to import or clear the foreign configuration. The import procedure and the clear procedure are described in [Section 3.7.3, "Importing or Clearing a Foreign Configuration."](#)

### 3.7.4 Migrating the RAID Level of a Virtual Drive

As the amount of data and the number of drives in your system increase, you can use RAID-level migration to change a virtual drive from one RAID level to another. You do not have to power down or reboot the system. When you migrate a virtual drive, you can keep the same number of drives, or you can add drives. You can use the WebBIOS CU to migrate the RAID level of an existing virtual drive.

**Note:** While you can apply RAID-level migration at any time, LSI recommends that you do so when there are no reboots. Many operating systems issues I/O operations serially (one at a time) during boot. With a RAID-level migration running, a boot can often take more than 15 minutes.

Migrations are allowed for the following RAID levels:

- RAID 0 to RAID 1
- RAID 0 to RAID 5
- RAID 0 to RAID 6
- RAID 1 to RAID 0
- RAID 1 to RAID 5
- RAID 1 to RAID 6
- RAID 5 to RAID 0
- RAID 5 to RAID 6
- RAID 6 to RAID 0
- RAID 6 to RAID 5

[Table 3.3](#) lists the number of additional drives required when you change the RAID level of a virtual drive.

**Table 3.3 Additional Drives Required for RAID-Level Migration**

From RAID Level to RAID Level	Original Number of Drives in Drive Group	Additional Drives Required
RAID 0 to RAID 1	RAID 0: 1 drive	1
RAID 0 to RAID 5	RAID 0: 1 drive	2
RAID 0 to RAID 5	RAID 0: 1 drive	3
RAID 1 to RAID 5	RAID 1: 2 drives	1
RAID 1 to RAID 6	RAID 1: 2 drives	1

Follow these steps to migrate the RAID level:

**Caution:** Back up any data that you want to keep before you change the RAID level of the virtual drive.

1. On the main WebBIOS CU screen, select a virtual drive.
2. Click **Virtual Drives**.
3. When the Virtual Drive screen appears, select **Migration only** (and skip to [step 6](#)) or **Migration with addition** in the right panel.
4. If you selected **Migration with addition**, select one or more drives from the small window in the lower right of the screen.
5. Select a new RAID level from the drop-down menu on the right. The available RAID levels are limited, based on the current RAID level of the virtual drive plus the number of drives available.
6. When you have made your selections, click **Go** at the bottom of the right panel.
7. When the message appears, confirm that you want to migrate the RAID level of the virtual drive.

A reconstruction operation begins on the virtual drive. You must wait until the reconstruction is completed before you perform any other tasks in the WebBIOS CU.

# Chapter 4

## MegaRAID Command Tool

---

The MegaRAID Command Tool (CT) is a command line interface (CLI) application for SAS. You can use this utility to configure, monitor, and maintain MegaRAID SAS RAID controllers and the devices connected to them.

**Note:** The CT supports only the MegaRAID controller. It supports SAS and SATA II, but it does not support other types of MegaRAID controllers, such as U320, SATA I, or IDE.

**Note:** The IA-64 release for Windows is similar to the 32-bit release, so you can follow the 32-bit instructions. 32-bit applications that were validated on an x64 system, such as the Intel Markette system, can use the 32-bit instructions, also.

This chapter has the following sections:

- [Section 4.1, “Product Overview”](#)
- [Section 4.2, “Novell NetWare, SCO, Solaris, FreeBSD, and DOS Operating System Support”](#)
- [Section 4.3, “Command Line Abbreviations and Conventions”](#)
- [Section 4.4, “Controller Property-Related Options”](#)
- [Section 4.5, “Patrol Read-Related Controller Properties”](#)
- [Section 4.6, “BIOS-Related Properties”](#)
- [Section 4.7, “Battery Backup Unit-Related Properties”](#)
- [Section 4.8, “Options for Displaying Logs Kept at Firmware Level”](#)
- [Section 4.9, “Configuration-Related Options”](#)
- [Section 4.10, “Virtual Drive-Related Options”](#)
- [Section 4.11, “Drive-Related Options”](#)

- [Section 4.12, “Enclosure-Related Options”](#)
  - [Section 4.13, “Flashing the Firmware”](#)
  - [Section 4.14, “SAS Topology”](#)
  - [Section 4.15, “Diagnostic-Related Options”](#)
  - [Section 4.16, “Miscellaneous Options”](#)
- 

## 4.1 Product Overview

The MegaCLI Configuration Utility is a command line interface application you can use to manage MegaRAID SAS RAID controllers. You can use MegaCLI Configuration Utility to perform the following tasks:

- Configure MegaRAID SAS RAID controllers and attached devices
- Display information about virtual drives and drives for the controller and other storage components
- Display ongoing progress for operations on drives and virtual drives
- Change properties for the virtual drives and drives for the controller and other storage components
- Set, retrieve, and verify controller default settings
- Change the firmware on the controllers
- Monitor the RAID storage systems
- Support RAID levels 0, 1, 5, 6, 10, 50, and 60 (depending on the RAID controller)
- Create and use scripts with the scriptable CLI tool
- Configure drive into groups and virtual drives on the controller
- Display configuration information for the controller, drives, and virtual drives
- Change virtual drive properties on the controller
- Change drive properties on the controller
- Display controller properties
- Load configuration to the controller from a file
- Save the controller configuration to a file

- Start or stop a rebuild, consistency check (CC), or initialization operation
- Enable or disable a background initialization (BGI)
- Stop or display an ongoing background initialization
- Start or display a reconstruction
- Start or stop patrol read
- Set and retrieve patrol read related settings
- Flash new firmware on the SAS RAID controller
- Read and program NVRAM and flash memory directly into DOS
- Display relevant messages on the console and/or in the log file
- Display controller data using one command
- Exit with predefined success or failure exit codes
- Scan, preview, and import foreign configurations
- Set predefined environment variables, such as the number of controllers and virtual drives
- Display firmware event logs
- Display help for how to use the command line options:
- Display battery unit properties
- Display enclosure properties
- Display and set connector mode on supported controllers

The following sections describe the command line options in the MegaCLI Configuration Utility that you can use to perform these functions.

Note : The MegaCLI Configuration Utility has support for the Intel® Itanium (64-bit) platform. MegaCLI is the only application currently supported on IPF system.

---

## 4.2 Novell NetWare, SCO, Solaris, FreeBSD, and DOS Operating System Support

The MegaCLI Configuration Utility functions under the Novell® NetWare®, SCO® OpenServer™, SCO UnixWare®, Solaris, FreeBSD, and DOS operating systems in the same way that it does under the Windows and Linux operating systems. All of the commands supported for the Windows and Linux operating systems are supported for the NetWare, SCO, and Solaris operating systems as well.

For the SCO OpenServer and SCO UnixWare operating systems, LSI provides an executable file that you can execute from any folder, and an image of the same executable file on a floppy drive. The image filename is `MegaCLI.image`. The floppy disk is provided so that you can distribute MegaCLI and install the executable file later as needed.

For the Solaris operating system, LSI provides an executable file that you can execute from any folder. No installation is required.

For the Novell NetWare operating system, LSI provides an executable file, `MegaCLI.nlm`, that you can execute from any folder. No installation is required. The output of all of the commands appears in the console window.

---

## 4.3 Command Line Abbreviations and Conventions

This section explains the abbreviations and conventions used with MegaCLI Configuration Utility commands.

### 4.3.1 Abbreviations Used in the Command Line

[Table 4.1](#) lists the abbreviations for the virtual drive parameters used in the following sections.

**Table 4.1 Command Line Abbreviations**

Abbreviation	Description
WB	WriteBack write policy
WT	WriteThrough write policy
ADRA	Adaptive Read Ahead read policy
RA	Read Ahead read policy
NORA	Normal Read policy (No read ahead)
DIO	Direct I/O cache policy
CIO	Cached I/O cache policy

### 4.3.2 Conventions

There are some options for which you can specify multiple values. You can enter commands for a single controller (-aN), multiple controllers (-a0,1,2) or work on all present controllers (-aALL). This is denoted as -aN|-a0,1,2|-aALL in this document and specifies that you can enter commands for one controller, multiple controllers, or all controllers.

**Note :** All options in the MegaRAID Command Tool are position-dependent, unless otherwise specified.

[Table 4.2](#) describes the conventions used in the options.

**Table 4.2 Conventions**

Convention	Description
	Specifies "or," meaning you can choose between options.
-aN	N specifies the controller number for the command.
-a0,1,2	Specifies the command is for controllers 0, 1, and 2. You can select two or more controllers in this manner.
-aALL	Specifies the command is for all controllers.
-Lx	x specifies the virtual drive number for the command.
-L0,1,2	Specifies the command is for virtual drives 0, 1, and 2. You can select two or more virtual drives in this manner.
-Lall	Specifies the command is for all virtual drives.

**Table 4.2 Conventions (Cont.)**

Convention	Description
[E0:S0,E1,S1,...]	<p>Specifies when one or more physical devices need(s) to be specified in the command line. Each [E:S] pair specifies one physical device where E means device ID of the enclosure in which a drive resides, and S means the slot number of the enclosure.</p> <p>In the case of a physical device directly connected to the SAS port on the controller, with no enclosure involved, the format of [:S] can be used where S means the port number on the controller. For devices attached through the backplane, the firmware provides an enclosure device ID and MegaCLI expects the user input in the format of [E:S]. In the following sections, only the format, [E:S], is used in the command descriptions, although both formats are valid.</p>
[ ]	<p>Indicates that the parameter is optional except when it is used to specify physical devices. For example, [WT] means the write policy (WriteThrough) is optional.</p> <p>If you enter WT at the command line, the application will use WriteThrough write policy for the virtual drive. Otherwise, it uses the default value for the parameter.</p>
{ }	<p>Indicates that the parameters are grouped and that they must be given at the same time.</p>

You can specify the `-Silent` command line option for all possible functions of the MegaCLI Configuration Utility. If you enter this option at the command line, no message displays on the screen.

---

## 4.4 Controller Property-Related Options

You can use the commands in this section to set or display properties related to the controller(s), such as the virtual drive parameters and factory defaults.

## 4.4.1 Display Controller Properties

Use the command in [Table 4.3](#) to display parameters for the selected controller(s).

**Table 4.3 Controller Parameters**

<b>Convention</b>	<code>MegaCli -AdpAllinfo -aN -a0,1,2 -aALL</code>
<b>Description</b>	Displays information about the controller, including cluster state, BIOS, alarm, firmware version, BIOS version, battery charge counter value, rebuild rate, bus number/device number, present RAM, memory size, serial number of the board, and SAS address.

## 4.4.2 Display Number of Controllers Supported

Use the command in [Table 4.3](#) to display the number of controllers supported on the system.

**Table 4.4 Number of Controllers Supported**

<b>Convention</b>	<code>MegaCli -AdpCount</code>
<b>Description</b>	Displays the number of controllers supported on the system and returns the number to the operating system.

## 4.4.3 Enable or Disable Automatic Rebuild

Use the command in [Table 4.5](#) to turn automatic rebuild on or off for the selected controller(s). If you have configured hot spares and enabled automatic rebuild, the RAID controller automatically tries to use them to rebuild failed drives. Automatic rebuild also controls whether a rebuild will start when a drive that was part of the drive group is reinserted.

**Table 4.5 Enable or Disable Automatic Rebuild**

<b>Convention</b>	<code>MegaCli -AdpAutoRbld -Enbl -Dsbl -Dsply -aN -a0,1,2 -aALL</code>
<b>Description</b>	Enables or disables automatic rebuild on the selected controller(s). The <code>-Dsply</code> option shows the status of the automatic rebuild state.

## 4.4.4 Flush Controller Cache

Use the command in [Table 4.6](#) to flush the controller cache on the selected controller(s). This option sends the contents of cache memory to the virtual drive(s). If the MegaRAID system must be powered down rapidly, you must flush the contents of the cache memory to preserve data integrity.

**Table 4.6 Cache Flush on Selected Controller**

<b>Convention</b>	MegaCli -AdpCacheFlush -aN -a0,1,2 -aALL
<b>Description</b>	Flushes the controller cache on the selected controller(s).

## 4.4.5 Set Controller Properties

This command sets the properties on the selected controller(s). For example, for {RebuildRate -val}, you can enter a percentage between 0 percent and 100 percent as the value for the rebuild rate. (The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives.) 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.

**Note:** LSI recommends the default rebuild rate of 30 percent, and the default patrol read rate of 30 percent.

Use the command in [Table 4.7](#) to display the list of properties you can set for the controller(s).

**Table 4.7 Set Controller Properties**

<b>Convention</b>	MegaCli -AdpSetProp {CacheFlushInterval -val} {RebuildRate -val}  {PatrolReadRate -val} {BgiRate -val} {CCRate -val}  {ReconRate -val}  {SpinupDriveCount -val} {SpinupDelay -val} {CoercionMode -val}   {ClusterEnable -val} {PredFailPollInterval -val}  {BatWarnDsbl -val}  {EccBucketSize -val} {EccBucketLeakRate -val} {AbortCConError - val} AlarmEnbl   AlarmDsbl   AlarmSilence  {SMARTCpyBkEnbl -val}   - AutoDetectBackPlaneDsbl   -CopyBackDsbl   -LoadBalanceMode   -aN - a0,1,2 -aALL
-------------------	--

**Table 4.7 Set Controller Properties (Cont.)**

<b>Description</b>	<p>Sets the properties on the selected controller(s). The possible settings are:</p> <p>CacheFlushInterval: Cache flush interval in seconds. Values: 0 to 255.</p> <p>RebuildRate: Rebuild rate. Values: 0 to 100.</p> <p>PatrolReadRate: Patrol read rate. Values: 0 to 100.</p> <p>BgiRate: Background initialization rate. Values: 0 to 100.</p> <p>CCRate: Consistency check rate. Values: 0 to 100.</p> <p>ReconRate: Reconstruction rate. Values: 0 to 100.</p> <p>SpinupDriveCount: Max number of drives to spin up at one time. Values: 0 to 255.</p> <p>SpinupDelay: Number of seconds to delay among spinup groups. Values: 0 to 255.</p> <p>CoercionMode: Drive capacity Coercion mode. Values: 0 - None, 1 - 128 Mbytes, 2 - 1 Gbytes.</p> <p>ClusterEnable: Cluster is enabled or disabled. Values: 0 - Disabled, 1 - Enabled.</p> <p>PredFailPollInterval: Number of seconds between predicted fail polls. Values: 0 to 65535.</p> <p>BatWarnDsbl: Disable warnings for missing battery or missing hardware. Values: 0 - Enabled, 1 - Disabled.</p> <p>EccBucketSize: Size of ECC single-bit-error bucket. Values: 0 to 255.</p> <p>EccBucketLeakRate: Leak rate (in minutes) of ECC single-bit-error bucket. Values: 0 to 65535.</p> <p>AbortCCOnError:</p> <p>AlarmEnbl: Set alarm to Enabled.</p> <p>AlarmDsbl: Set alarm to Disabled.</p> <p>AlarmSilence: Silence an active alarm.</p> <p>SMARTCpyBkEnbl: Enable copyback operation on Self-Monitoring Analysis and Reporting Technology (SMART) errors. Copyback is initiated when the first SMART error occurs on a drive that is part of a virtual drive.</p> <p>AutoDetectBackPlaneDsbl: Detect automatically if the backplane has been disabled.</p> <p>CopyBackDsbl: Disable or enable the copyback operation.</p> <p>LoadBalanceMode: Disable or enable the load balancing mode.</p>
--------------------	--

#### 4.4.6 Display Specified Controller Properties

Use the command in [Table 4.8](#) to display specified properties on the selected controller(s).

**Table 4.8 Display Specified Controller Properties**

<b>Convention</b>	<pre>MegaCli -AdpGetProp CacheFlushInterval   RebuildRate   PatrolReadRate   BgiRate   CCRate   ReconRate   SpinupDriveCount   SpinupDelay   CoercionMode   PredFailPollInterval   ClusterEnable   BatWarnDsbl   EccBucketSize   EccBucketLeakRate   EccBucketCount   AlarmDsply -a  -a0,1,2 -aALL</pre>
<b>Description</b>	<p>Displays the properties on the selected controller(s).</p> <p>EccBucketCount: Count of single-bit ECC errors currently in the bucket.</p> <p>See <a href="#">Table 4.7</a> for explanations of the other options.</p>

## 4.4.7 Set Factory Defaults

Use the command in [Table 4.9](#) to set the factory defaults on the selected controller(s).

**Table 4.9 Set Factory Defaults**

<b>Convention</b>	MegaCli -AdpFacDefSet -aN -a0,1,2 -aALL
<b>Description</b>	Sets the factory defaults on the selected controller(s).

## 4.4.8 Set SAS Address

Use the command in [Table 4.10](#) to set the SAS address on the selected controller(s).

**Table 4.10 Set SAS Address on Controller**

<b>Convention</b>	MegaCli -AdpSetSASA str[0-64] -aN
<b>Description</b>	Sets the controllers SAS address. This string must be a 64-digit hexadecimal number.

## 4.4.9 Set Time and Date on Controller

Use the command in [Table 4.11](#) to set the time and date on the selected controller(s).

**Table 4.11 Set Time and Date on Controller**

<b>Convention</b>	MegaCli -AdpSetTime yyyyymmdd HH:mm:ss -aN -a0,1,2 -aALL
<b>Description</b>	Sets the time and date on the controller. This command uses a 24-hour format. For example, 7 p.m. displays as 19:00:00. The order of date and time is reversible.

## 4.4.10 Display Time and Date on Controller

Use the command in [Table 4.12](#) to display the time and date on the selected controller(s).

**Table 4.12 Display Time and Date on Controller**

<b>Convention</b>	MegaCli -AdpGetTime -aN
<b>Description</b>	Displays the time and date on the controller. This command uses a 24-hour format. For example, 7 p.m. would display as 19:00:00.

---

## 4.5 Patrol Read-Related Controller Properties

You can use the commands in this section to select the settings for Patrol Read. A Patrol Read scans the system for possible drive errors that could lead to drive failure, then takes action to correct the errors. The goal is to protect data integrity by detecting drive failure before the failure can damage data. The corrective actions depend on the virtual drive configuration and the type of errors. Patrol Read affects performance; the more iterations there are, the greater the impact.

### 4.5.1 Set Patrol Read Options

Use the command in [Table 4.13](#) on the selected controller(s) to set the Patrol Read options.

**Table 4.13 Set Patrol Read Options**

<b>Convention</b>	<code>MegaCli -AdpPR -DsbL   EnblAuto   EnblMan   Start   Stop   Info -aN   -a0,1,2   -aALL</code>
<b>Description</b>	<p>Sets Patrol Read options on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"><li>-DsbL: Disables Patrol Read for the selected controller(s).</li><li>-EnblAuto: Enables Patrol Read automatically for the selected controller(s). This means Patrol Read will start automatically after the controller initialization is complete.</li><li>-EnblMan: Enables Patrol Read manually for the selected controller(s). This means that Patrol Read does not start automatically; it has to be started manually by selecting the <code>Start</code> command.</li><li>-Start: Starts Patrol Read for the selected controller(s).</li><li>-Stop: Stops Patrol Read for the selected controller(s).</li><li>-Info: Displays the following Patrol Read information for the selected controller(s):<ul style="list-style-type: none"><li>• Patrol Read operation mode</li><li>• Patrol Read execution delay value</li><li>• Patrol Read status</li></ul></li></ul>

---

## 4.5.2 Set Patrol Read Delay Interval

Use the command in [Table 4.14](#) on the selected controller(s) to set the time between Patrol Read iterations.

**Table 4.14 Set Patrol Read Delay Interval**

---

<b>Convention</b>	MegaCli -AdpPRSetDelay -Val -aN -a0,1,2 -aALL
<b>Description</b>	Sets the time between Patrol Read iterations on a single controller, multiple controllers, or all controllers: -Val: Sets delay time between Patrol Read iterations. The value is time of delay in hours. A value of zero means no delay and an immediate restart.

---

## 4.6 BIOS-Related Properties

You can use the commands in this section to select the settings for BIOS-related options.

### 4.6.1 Set or Display Bootable Virtual Drive ID

Use the command in [Table 4.15](#) to set or display the ID of the bootable virtual drive.

**Note:** This option does not write a boot sector to the virtual drive. The operating system will not load if the boot sector is incorrect.

**Table 4.15 Bootable Virtual Drive ID**

---

<b>Convention</b>	MegaCli -AdpBootDrive {-Set -Lx}   -Get -aN -a0,1,2 -aALL
<b>Description</b>	Sets or displays the bootable virtual drive ID: -Set: Sets the virtual drive as bootable so that during the next reboot, the BIOS will look for a boot sector in the specified virtual drive. -Get: Displays the bootable virtual drive ID.

---

## 4.6.2 Select BIOS Status Options

Use the command in [Table 4.16](#) to set the options for the BIOS status.

**Table 4.16 Options for BIOS Status**

<b>Convention</b>	MegaCli -AdpBIOS -Enbl -Dsbl -Dsply  SOE   BE -aN -a0,1,2 -aALL
<b>Description</b>	Sets BIOS options. The following are the settings you can select on a single controller, multiple controllers, or all controllers: -Enbl, -Dsbl, -Dsply: Enables, disables or displays the BIOS status on selected controller(s). -SOE: Stops on BIOS errors during POST for selected controller(s). When set to -SOE, 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. This is available only when you enable the BIOS status. -BE: Bypasses BIOS errors during POST. This is available only when you enable the BIOS status.

## 4.7 Battery Backup Unit-Related Properties

You can use the commands in this section to select the settings for BBU-related options.

### 4.7.1 Display BBU Information

Use the command in [Table 4.17](#) to display complete information about the BBU for the selected controller(s).

**Table 4.17 Display BBU Information**

<b>Convention</b>	MegaCli -AdpBbuCmd -aN -a0,1,2 -aALL
<b>Description</b>	Displays complete information about the BBU, such as status, capacity information, design information, and properties.

## 4.7.2 Display BBU Status Information

Use the command in [Table 4.18](#) to display complete information about the status of the BBU, such as temperature and voltage, for the selected controller(s).

**Table 4.18 Display BBU Status Information**

<b>Convention</b>	MegaCli -AdpBbuCmd -GetBbuStatus -aN -a0,1,2 -aALL
<b>Description</b>	<p>Displays complete information about the BBU status, such as the temperature and voltage. The information displays in the following formats:</p> <p><b>BBU Status for Adapter: xx</b> Battery Type: XXXXXX(string) Voltage: xx mV Current: xx mA Temperature: xx C° Firmware Status: xx Battery state: xx</p> <p><b>Gas Gauge Status:</b> Fully Discharged: Yes/No Fully Charged: Yes/No Discharging: Yes/No Initialized: Yes/No Remaining Time Alarm: Yes/No Remaining Capacity Alarm: Yes/No Discharge Terminated: Yes/No Over Temperature: Yes/No Charging Terminated: Yes/No Over Charged: Yes/No</p> <p>Additional status information displays differently for iBBU™ and BBU.</p> <p><b>For iBBU:</b> Relative State of Charge: xx Charger System State: xx Charger System Ctrl: xx Charging Current: xx mA Absolute State of Charge: xx% Max Error: xx%</p> <p><b>For BBU:</b> Relative State of Charge: xx Charger Status: xx Remaining Capacity: xx mAh Full Charge Capacity: mAh isSOHGood: Yes/No</p>

### 4.7.3 Display BBU Capacity

Use the command in [Table 4.19](#) to display the BBU capacity for the selected controller(s).

**Table 4.19 Display BBU Capacity Information**

---

<b>Convention</b>	MegaCli -AdpBbuCmd -GetBbuCapacityInfo -aN -a0,1,2 -aALL
<b>Description</b>	Displays BBU capacity information. The information displays in the following format: BBU Capacity Info for Adapter: x Relative State of Charge: xx% Absolute State of Charge: xx% Remaining Capacity: xx mAh Full Charge Capacity: xx mAh Run Time to Empty: xxx Min Average Time to Empty: xxx Min Average Time to Full: xxx Min Cycle Count: xx Max Error: xx%

---

### 4.7.4 Display BBU Design Parameters

Use the command in [Table 4.20](#) to display BBU design parameters for the selected controller(s).

**Table 4.20 Display BBU Design Parameters**

---

<b>Convention</b>	MegaCli -AdpBbuCmd -GetBbuDesignInfo -aN -a0,1,2 -aALL
<b>Description</b>	Displays information about the BBU design parameters. The information displays in the following formats: BBU Design Info for Adapter: x Date of Manufacture: mm/dd, yyyy Design Capacity: xxx mAh Design Voltage: mV Serial Number: 0xhhhh Pack Stat Configuration: 0xhhhh Manufacture Name: XXXXXX(String) Device Name: XXXXXX(String) Device Chemistry: XXXXXX(String)

---

## 4.7.5 Display Current BBU Properties

Use the command in [Table 4.21](#) to display the current BBU properties for the selected controller(s).

**Table 4.21 Display Current BBU Properties**

<b>Convention</b>	MegaCli -AdpBbuCmd -GetBbuProperties -aN -a0,1,2 -aALL
<b>Description</b>	Displays current properties of the BBU. The information displays in the following formats: BBU Properties for Adapter: x Auto Learn Period: xxx Sec Next Learn Time: xxxx Sec Learn Delay Interval: xx Hours Auto-Learn Mode: Warn via Event/Disabled/Enabled

## 4.7.6 Start BBU Learning Cycle

Use the command in [Table 4.22](#) to start the BBU learning cycle on the selected controller(s). A learning cycle is a battery calibration operation performed by the controller periodically (approximately every three months) to determine the condition of the battery.

**Table 4.22 Start BBU Learning Cycle**

<b>Convention</b>	MegaCli -AdpBbuCmd -BbuLearn -aN -a0,1,2 -aALL
<b>Description</b>	Starts the learning cycle on the BBU. No parameter is needed for this option.

## 4.7.7 Place Battery in Low-Power Storage Mode

Use the command in [Table 4.23](#) to place the battery into Low-Power Storage mode on the selected controller(s). This saves battery power consumption.

**Table 4.23 Place Battery in Low-Power Storage Mode**

<b>Convention</b>	MegaCli -AdpBbuCmd -BbuMfgSleep -aN -a0,1,2 -aALL
<b>Description</b>	Places the battery in Low-Power Storage mode. The battery automatically exits this state after 5 seconds.

## 4.7.8 Set BBU Properties

Use the command in [Table 4.24](#) to set the BBU properties on the selected controller(s) after reading from the file.

**Table 4.24 Set BBU Properties**

<b>Convention</b>	MegaCli -AdpBbuCmd -SetBbuProperties -f<fileName> -aN -a0,1,2 -aALL
<b>Description</b>	Sets the BBU properties on the selected controller(s) after reading from the file. The information displays in the following formats: autoLearnPeriod = 1800Sec nextLearnTime = 12345678Sec Seconds past 1/1/2000 learnDelayInterval = 24hours Not greater than 7 days autoLearnMode = 0 0 – Enabled, 1 - Disabled, 2 – WarnViaEvent.  1. NOTE: You can change only two of these parameters, learnDelayInterval and autoLearnMode.

## 4.8 Options for Displaying Logs Kept at Firmware Level

Use the commands in this section to select the display settings for the event log and BBU terminal log, which are kept at the firmware level.

### 4.8.1 Event Log Management

Use the command in [Table 4.25](#) to manage the event entries in the event log for the selected controller(s).

**Table 4.25 Event Log Management**

<b>Convention</b>	MegaCli -AdpEventLog -GetEventlogInfo   -GetEvents   GetSinceShutdown   GetSinceReboot   IncludeDeleted   {GetLatest <number>} -f <filename>   Clear -aN -a0,1,2 -aALL   {GetCCIncon} -f <filename> -LX -L0,2,5... -LALL -aN -a0,1,2 -aALL
-------------------	--

**Table 4.25 Event Log Management (Cont.)**

<b>Description</b>	<p>Manages event log entries. The following are the settings you can select on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"> <li>-GetEventlogInfo: Displays overall event information such as total number of events, newest sequence number, oldest sequence number, shutdown sequence number, reboot sequence number, and clear sequence number.</li> <li>-GetEvents: Gets event log entry details. The information shown consists of total number of entries available at firmware side since the last clear and details of each entries of the error log. <code>Start_entry</code> specifies the initial event log entry when displaying the log.</li> <li>-GetSinceShutdown: Displays all of the events since last controller shutdown.</li> <li>-GetSinceReboot: Displays all of the events since last controller reboot.</li> <li>-IncludeDeleted: Displays all events, including deleted events.</li> <li>-GetLatest: Displays the latest number of events, if any exist. The event data will be writtend to the file in reverse order.</li> <li>-Clear: Clears the event log for the selected controller(s).</li> <li>-GetCCIncon:</li> </ul>
--------------------	--

## 4.8.2 Set BBU Terminal Logging

Use the command in [Table 4.26](#) to set the BBU terminal logging for the selected controller(s).

**Table 4.26 Set BBU Terminal Logging**

<b>Convention</b>	<pre>MegaCli -FwTermLog -Bbuoff   -BbuoffTemp   -Bbuon   -BbuGet   -Dsply   -Clear -aN   -a0,1,2   -aALL</pre>
<b>Description</b>	<p>Sets BBU terminal logging options. The following are the settings you can select on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"> <li>-Bbuoff: Turns off the BBU for firmware terminal logging. To turn off the BBU for logging, you have to shut down your system or turn off the power to the system after you run the command.</li> <li>-BbuoffTemp: Temporarily turns off the BBU for TTY (firmware terminal) logging. The battery will be turned on at the next reboot.</li> <li>-Bbuon: Turns on the BBU for TTY (firmware terminal) logging.</li> <li>-BbuGet: Displays the current BBU settings for TTY logging.</li> <li>-Dsply: Displays the TTY log (firmware terminal log) entries with details on the given controllers. The information shown consists of the total number of entries available at a firmware side.</li> <li>-Clear: Clears the TTY log.</li> </ul>

---

## 4.9 Configuration-Related Options

You can specify the drives by using the Enclosure ID:Slot ID for SAS controllers. This assumes that all drives are connected to the controller through an enclosure. If the drives are not connected to an enclosure, it is assumed that they are connected to Enclosure 0. In this case there is no slot, so you can use the `pdlst` command to get the slot equivalent number. (This applies to all commands that use the Enclosure ID:Slot ID format.) MegaCLI expects the input in `[:S]` format for directly attached devices.

In the following options, `[E0:S0, E1:S1]` specifies the enclosure ID and slot ID for the drive.

### 4.9.1 Create a RAID Drive Group from All Unconfigured Good Drives

Use the command in [Table 4.28](#) to create one RAID drive group out of all of the unconfigured good drives, and a hot spare, if desired. This is for RAID levels 0, 5, 6, 10, 50, or 60. All free drives are used to create a new drive group and, if desired, one hot spare drive. If it is not possible to use all of the free drives, the command will abort with a related error level. If there are drives of different capacities, the largest drive is used to make the hot spare.

**Note:** A virtual drive cannot have both SAS drives and SATA drives. Therefore, if the remaining free drives are SAS and SATA, a drive group cannot be created. The command will abort with a related error level.

**Note:** Firmware supports only 32 drives per drive group, so if there are more than 32 unconfigured good drives, MegaCLI cannot configure any of the drives, and the command will abort.

**Table 4.27 Create a Drive Group from All of the Unconfigured Drives**

---

<b>Convention</b>	<code>MegaCli -CfgLDAdd -R0 -R1 -R5 -R6[E0:S0,E1:S1,...] [WT   WB] [NORA   RA   ADRA] [Direct   Cached] [CachedBadBBU NoCachedBadBBU] [ GHSP   DHSP] -spn x -ax   -Force</code>
-------------------	---

---

**Table 4.27 Create a Drive Group from All of the Unconfigured Drives (Cont.)**

---

<b>Description</b>	<p>Creates one RAID drive group out of all of the unconfigured good drives, and a hot spare, if desired. This is for RAID levels 0, 5, 6, 10, 50, or 60. All free drives are used to create a new drive group and, if desired, one hot spare drive.</p> <p>-Rx[E0:S0,...]: Specifies the RAID level and the drive enclosure/slot numbers used to construct a drive group.</p> <p>-WT (Write through), WB (Write back): Selects write policy.</p> <p>-NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead): Selects read policy.</p> <p>-Direct, -Cached: Selects cache policy.</p> <p>[{CachedBadBBU NoCachedBadBBU }]: Specifies whether to use write cache when the BBU is bad.</p> <p>GHSP, DHSP: Specifies that the hot spare is global or dedicated.</p> <p>-spn: Specifies the maximum number of spans that a spanned drive group can have.</p> <p>-Force: Specifies that drive coercion is used to make the capacity of the drives compatible. Drive coercion is a tool for forcing drives of varying capacities to the same capacity so they can be used in a drive group.</p>
--------------------	---

---

## 4.9.2 Add RAID 0, 1, 5, or 6 Configuration

Use the command in [Table 4.28](#) to add a RAID level 0, 1, 5, or 6 configuration to the existing configuration on the selected controller. For RAID levels 10, 50, or 60, see [Section 4.9.3, “Add RAID 10, 50, or 60 Configuration.”](#)

**Table 4.28 Add RAID 0, 1, 5, or 6 Configuration**

---

<b>Convention</b>	<pre>MegaCli -CfgLDAdd -R0 -R1 -R5 -R6[E0:S0,E1:S1,...] [WT   WB] [NORA   RA   ADRA] [Direct   Cached] [CachedBadBBU NoCachedBadBBU] [- szXXXXXXXX [-szYYYYYYY [... ]]] [-strpszM] [-Hsp[E5:S5,...]] [- afterLdX] -aN</pre>
-------------------	---

---

**Table 4.28 Add RAID 0, 1, 5, or 6 Configuration (Cont.)**

---

<b>Description</b>	<p>Adds a RAID level 0, 1, 5, or 6 configuration to a specified controller. Even if no configuration is present, you have the option to write the configuration to the controller.</p> <p>Note that RAID 1 supports up to 32 drives in a single span of 16 drive groups. RAID 1 requires an even number of drives, as data from one drive is mirrored to the other drive in each RAID 1 drive group.</p> <p>-Rx[E0:S0, . . .]: Specifies the RAID level and the drive enclosure/slot numbers to construct a drive group.</p> <p>-WT (Write through), WB (Write back): Selects write policy.</p> <p>-NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead): Selects read policy.</p> <p>-Cached, -Direct: Selects cache policy.</p> <p>[{CachedBadBBU NoCachedBadBBU }]: Specifies whether to use write cache when the BBU is bad.</p> <p>-szXXXXXXXX: Specifies the capacity for the virtual drive, where XXXX is a decimal number of Mbytes. However, the actual capacity of the virtual drive can be smaller, because the driver requires the number of blocks from the drives in each virtual drive to be aligned to the stripe size. If multiple size options are specified, CT configures the virtual drives in the order of the options entered in the command line. The configuration of a particular virtual drive will fail if the remaining capacity of the drive group is too small to configure the virtual drive with the specified capacity. This option can also be used to create a configuration on the free space available in the drive group.</p> <p>-strpszM: Specifies the stripe size, where the stripe size values are 8, 16, 32, 64, 128, 256, 512, or 1024 KBytes.</p> <p>Hsp[E5:S5, . . .]: Creates hot spares when you create the configuration. The new hot spares will be dedicated to the virtual drive used in creating the configuration. This option does not allow you to create global hot spares. To create global hot spares, you must use the -PdHsp command with proper subcommands. You can also use this option to create a configuration on the free space available in the virtual drive. You can specify which free slot should be used by specifying the -AfterLdX. This command is optional. By default, the application uses the first free slot available in the virtual drive. This option is valid only if the virtual drive is already used for configuration.</p>
--------------------	--

---

### 4.9.3 Add RAID 10, 50, or 60 Configuration

Use the command in [Table 4.29](#) to add a RAID 10, RAID 50, or RAID 60 configuration to the existing configuration on the selected controller. For RAID levels 0, 1, 5, or 6, see [Section 4.9.2, “Add RAID 0, 1, 5, or 6 Configuration.”](#)

**Table 4.29 Add RAID 10, 50, or 60 Configuration**

<b>Convention</b>	<code>MegaCli -CfgSpanAdd -R10 -R50 R60 -Array0[E0:S0,E1:S1,...] -Array1[E0:S0,E1:S1,...] [...] [WT   WB] [NORA   RA   ADRA] [Direct   Cached] [CachedBadBBU NoCachedBadBBU] [-szXXXXXXXX [-szYYYYYYY [...]]] [-strpszM] [-afterLdX] -aN -a0,1,2 -aALL</code>
<b>Description</b>	<p>Creates a RAID level 10, 50, or 60 (spanned) configuration from the specified drive groups. Even if no configuration is present, you must use this option to write the configuration to the controller.</p> <p>Note that RAID 10 supports up to eight spans with a maximum of 32 drives in each span. (There are factors, such as the type of controller, that limit the number of drives you can use.) RAID 10 requires an even number of drives, as data from one drive is mirrored to the other drive in each RAID 1 drive group. You can have an even or odd number of spans.</p> <p>Multiple drive groups are specified using the <code>-ArrayX[E0:S0,...]</code> option. (Note that <i>X</i> starts from 0, not 1.) All of the drive groups must have the same number of drives. At least two drive groups must be provided. The order of options {WT  WB} {NORA   RA   ADRA} {Direct   Cached} is flexible.</p> <p>The size option, <code>-szXXXXXXXX</code>, can be accepted to allow slicing in the spanned drive groups if the controller supports this feature. The <code>[-afterLdX]</code> option is accepted if the size option is accepted. CT exits and does not create a configuration if the size or the afterLd option is specified but the controller does not support slicing in the spanned drive groups.</p>

### 4.9.4 Clear the Existing Configuration

Use the command in [Table 4.30](#) to clear the existing configuration on the selected controller(s).

**Table 4.30 Clear Existing Configuration**

<b>Convention</b>	<code>MegaCli -CfgClr -aN -a0,1,2 -aALL</code>
<b>Description</b>	Clears the existing configuration.

## 4.9.5 Save the Configuration on the Controller

Use the command in [Table 4.31](#) to save the configuration for the selected controller(s) to the given filename.

**Table 4.31 Save Configuration on the Controller**

<b>Convention</b>	MegaCli -CfgSave -f FileName -aN
<b>Description</b>	Saves the configuration for the selected controller(s) to the given filename.

## 4.9.6 Restore the Configuration Data from File

Use the command in [Table 4.32](#) to read the configuration from the file and load it on the selected controller(s). You can restore the read/write properties and RAID configuration using hot spares.

**Table 4.32 Restore Configuration Data from File**

<b>Convention</b>	MegaCli -CfgRestore -f FileName -aN
<b>Description</b>	Reads the configuration from the file and loads it on the controller. MegaCLI can store or restore all read and write controller properties, all read and write properties for virtual drives, and the RAID configuration including hot spares. Note the following: <ul style="list-style-type: none"><li>• MegaCLI does not validate the setup when restoring the RAID configuration.</li><li>• The -CfgSave option stores the configuration data and controller properties in the file. Configuration data has only the device ID and sequence number information of the drives used in the configuration. The CfgRestore option will fail if the same device IDs of the drives are not present.</li></ul>

## 4.9.7 Manage Foreign Configuration Information

Use the command in [Table 4.33](#) to manage configurations from other controllers, called *foreign configurations*, for the selected controller(s). You can scan, preview, import, and clear foreign configurations.

**Note:** The actual status of virtual drives and drives can differ from the information displayed in the -Scan option. LSI suggests that you run -Preview before you import a foreign configuration.

**Table 4.33 Manage Foreign Configuration Information**

<b>Convention</b>	MegaCli -CfgForeign -Scan   {-Preview   -Dsply  -Import   -Clear [FID]} -aN -a0,1,2 -aALL
<b>Description</b>	Manages foreign configurations. The options for this command are: -Scan: Scans and displays available foreign configurations. -Preview: Provides a preview of the imported foreign configuration. The foreign configuration ID (FID) is optional. -Dsply: Displays the foreign configuration. -Import: Imports the foreign configuration. The FID is optional. -Clear [FID]: Clears the foreign configuration. The FID is optional.

## 4.9.8 Delete Specified Virtual Drive(s)

Use the command in [Table 4.34](#) to delete one, multiple, or all virtual drives on the selected controller(s).

**Table 4.34 Delete Specified Virtual Drives**

<b>Convention</b>	MegaCli -CfgLDDel -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
<b>Description</b>	Deletes the specified virtual drive(s) on the selected controller(s). You can delete one virtual drive, multiple virtual drives, or all of the selected virtual drives on selected controller(s).

## 4.9.9 Display the Free Space

Use the command in [Table 4.35](#) to display the free space that is available to use for configuration on the selected controller(s).

**Table 4.35 Display Free Space**

<b>Convention</b>	MegaCli -CfgFreeSpaceInfo -aN -a0,1,2 -aALL
<b>Description</b>	Displays all of the free space available for configuration on the selected controller(s). The information displayed includes the number of drive groups, the number of spans in each drive group, the number of free space slots in each drive group, the start block, and the size (in both blocks and megabytes) of each free space slot.

---

## 4.10 Virtual Drive-Related Options

You can use the commands in this section to select settings for the virtual drives and perform actions on them.

### 4.10.1 Display Virtual Drive Information

Use the command in [Table 4.36](#) to display virtual drive information for the selected controller(s).

**Table 4.36 Display Virtual Drive Information**

<b>Convention</b>	<code>MegaCli -LDInfo -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL</code>
<b>Description</b>	Displays information about the virtual drive(s) on the selected controller(s). This information includes the name, RAID level, RAID level qualifier, capacity in megabytes, state, stripe size, number of drives, span depth, cache policy, access policy, and ongoing activity progress, if any, including initialization, background initialization, consistency check, and reconstruction.

### 4.10.2 Change the Virtual Drive Cache and Access Parameters

Use the command in [Table 4.37](#) to change the cache policy and access policy for the virtual drive(s) on the selected controller(s).

**Table 4.37 Change Virtual Drive Cache and Access Parameters**

<b>Convention</b>	<code>MegaCli -LDSetProp WT   WB NORA  RA   ADRA -Cached Direct  CachedBadBBU NoCachedBadBBU}   -RW RO Blocked   {-Name nameString}   -EnDskCache DisDskCache -Lx  -L0,1,2 -Lall -aN -a0,1,2 -aALL</code>
<b>Description</b>	Allows you to change the following virtual drive parameters: -WT (Write through), WB (Write back): Selects write policy. -NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead): Selects read policy. -Cached, -Direct: Selects cache policy. -CachedBadBBU NoCachedBadBBU : Specifies whether to use write cache when the BBU is bad. -RW, -RO, Blocked: Selects access policy. -EnDskCache: Enables drive cache. -DisDskCache: Disables drive cache.

### 4.10.3 Display the Virtual Drive Cache and Access Parameters

Use the command in [Table 4.38](#) to display cache and access parameters for the virtual drive(s) on the selected controller(s).

**Table 4.38 Display Virtual Drive Cache and Access Parameters**

<b>Convention</b>	MegaCli -LDGetProp -Cache   -Access   -Name   -DskCache -Lx -L0,1,2  -Lall -aN -a0,1,2 -aALL
<b>Description</b>	Displays the cache and access policies of the virtual drive(s): -Cache: -Cached, Direct: Displays cache policy. -WT (Write through), WB (Write back): Selects write policy. -NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead): Selects read policy. -Access: -RW, -RO, Blocked: Displays access policy. -DskCache: Displays drive cache policy.

### 4.10.4 Manage Virtual Drives Initialization

Use the command in [Table 4.39](#) to manage initialization of the virtual drive(s) on the selected controller(s).

**Table 4.39 Manage Virtual Drive Initialization**

<b>Convention</b>	MegaCli -LDInit {-Start [Fast   Full]}   -Abort -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
<b>Description</b>	Allows you to select the following actions for virtual drive initialization: -Start: Starts the initialization (writing 0s) on the virtual drive(s) and displays the progress (this is optional). The fast initialization option initializes the first and last 8 Mbyte areas on the virtual drive. The full option allows you to initialize the entire virtual drive. -Abort: Aborts the ongoing initialization on the virtual drive(s). -ShowProg: Displays the snapshot of the ongoing initialization, if any. -ProgDsply: Displays the progress of the ongoing initialization. The routine continues to display the progress until at least one initialization is completed or a key is pressed.

## 4.10.5 Manage a Consistency Check

Use the command in [Table 4.40](#) to manage a data consistency check (CC) on the virtual drives for the selected controller(s).

**Table 4.40 Manage Consistency Check**

<b>Convention</b>	MegaCli -LDCC -Start -Abort -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
<b>Description</b>	Allows you to select the following actions for a data CC: -Start: Starts a CC on the virtual drive(s), then displays the progress (optional) and time remaining. -Abort: Aborts an ongoing CC on the virtual drive(s). -ShowProg: Displays a snapshot of an ongoing CC. -ProgDsply: Displays ongoing CC progress. The progress displays until at least one CC is completed or a key is pressed.

## 4.10.6 Manage a Background Initialization

Use the command in [Table 4.41](#) to enable, disable, or suspend background initialization (BGI), as well as display initialization progress on the selected controller(s).

**Table 4.41 Manage Background Initialization**

<b>Convention</b>	MegaCli -LDBI -Enbl -Dsbl GetSetting -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
<b>Description</b>	Manages background initialization options. The following are the background initialization settings you can select on a single controller, multiple controllers, or all controllers: -Enbl, -Dsbl: Enables or disables the background initialization on the selected controller(s). -ProgDsply: Displays an ongoing background initialization in a loop. This function completes only when all background initialization processes complete or you press a key to exit. -ShowProg: Displays the current progress value. - GetSetting: Displays current background initialization setting ( <i>Enabled</i> or <i>Disabled</i> ).

## 4.10.7 Perform a Virtual Drive Reconstruction

Use the command in [Table 4.42](#) to perform a reconstruction of the virtual drive(s) on the selected controller(s).

**Table 4.42 Virtual Drive Reconstruction**

<b>Convention</b>	<code>MegaCli -LDRcon {-Start -Rx [Add   Rmv PhysDrv[E0:S0,E1:S1,...] ] }  -ShowProg -ProgDsply -Lx -aN</code>
<b>Description</b>	Controls and manages virtual drive reconstruction. The following are the virtual drive reconstruction settings you can select on a single controller: -Start: Starts a reconstruction of the selected virtual drive to a new RAID level. -Rx: Changes the RAID level of the virtual drive when you start reconstruction. You might need to add or remove a drive to make this possible. -Start -Add PhysDrv[E0:S0,E1:S1...]: Adds listed drives to the virtual drive and starts reconstruction on the selected virtual drive. -Start -Rmv PhysDrv[E0:S0,E1:S1...]: Removes one drive from the existing virtual drives and starts a reconstruction. -ShowProg: Displays a snapshot of the ongoing reconstruction process. -ProgDsply: Allows you to view the ongoing reconstruction. The routine continues to display progress until at least one reconstruction is completed or a key is pressed.

## 4.10.8 Display Information about Virtual Drives and Drives

Use the command in [Table 4.43](#) to display information about the virtual drives and drives for the selected controller(s), such as the number of virtual drives, RAID level, and drive capacity.

**Table 4.43 Display Virtual Drive and Drive Information**

<b>Convention</b>	<code>MegaCli -LDPDInfo -aN -a0,1,2 -aALL</code>
<b>Description</b>	Displays information about the present virtual drive(s) and drive(s) on the selected controller(s). Displays information including the number of virtual drives, the RAID level of the virtual drives, and drive capacity information, which includes raw capacity, coerced capacity, uncoerced capacity, and the SAS address.

## 4.10.9 Display the Number of Virtual Drives

Use the command in [Table 4.44](#) to display the number of virtual drives attached to the controller.

**Table 4.44 Display Number of Virtual Drives**

---

<b>Convention</b>	MegaCli -LDGetNum -aN -a0,1,2 -aALL
<b>Description</b>	Displays the number of virtual drives attached to the controller. The return value is the number of virtual drives.

---

## 4.11 Drive-Related Options

You can use the commands in this section to select settings for the drives and perform actions on them.

### 4.11.1 Display Drive Information

Use the command in [Table 4.45](#) to display information about the drives on the selected controller(s).

**Table 4.45 Display Drive Information**

---

<b>Convention</b>	MegaCli -PDInfo -PhysDrv[E0:S0,E1:S1...] -aN -a0,1,2 -aALL
<b>Description</b>	Provides information about the drives connected to the enclosure and controller slot. This includes information such as the enclosure number, slot number, device ID, sequence number, drive type, capacity (if a drive), foreign state, firmware state, and inquiry data.  For SAS devices, this includes additional information such as the SAS address of the drive. For SAS expanders, this includes additional information such as the number of devices connected to the expander.

---

## 4.11.2 Set the Drive State to Online

Use the command in [Table 4.46](#) to set the state of a drive to *Online*. In an online state, the drive is working normally and is a part of a configured virtual drive.

**Table 4.46 Set Drive State to Online**

<b>Convention</b>	MegaCli -PDOnline -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Changes the drive state to <i>Online</i> .

## 4.11.3 Set the Drive State to Offline

Use the command in [Table 4.47](#) to set the state of a drive to *Offline*. In the offline state, the virtual drive is not available to the RAID controller.

**Table 4.47 Set Drive State to Offline**

<b>Convention</b>	MegaCli -PDOffline -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Changes the drive state to <i>Offline</i> .

## 4.11.4 Change the Drive State to Unconfigured Good

Use the command in [Table 4.48](#) to change the state of a drive from *Unconfigured-Bad* to *Unconfigured-Good*.

**Table 4.48 Change Drive State to Unconfigured Good**

<b>Convention</b>	MegaCli -PDMakeGood -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Changes the drive state to <i>Offline</i> .

## 4.11.5 Change Drive State

Use the command in [Table 4.49](#) to change the drive state, as it relates to hot spares, and to associate the drive to an enclosure and virtual drive for the selected controller(s).

**Table 4.49 Change Drive State**

<b>Convention</b>	<code>MegaCli -PDHSP {-Set [{-Dedicated -ArrayN  -Array0,1...}] [-EnclAffinity] [-nonRevertible] }   -Rmv -PhysDrv[E0:S0,E1:S1,...] -aN -a0,1,2 -aALL</code>
<b>Description</b>	Changes the drive state (as it relates to hot spares) and associates the drive to an enclosure and virtual drive on a single controller, multiple controllers, or all controllers: -Set: Changes the drive state to <i>dedicated hot spare</i> for the enclosure. -Rmv: Changes the drive state to <i>ready</i> (removes the hot spare). -EnclAffinity: Associates the hot spare to a selected enclosure. -Array0: Dedicates the hot spare to a specific virtual drive.

## 4.11.6 Manage a Drive Initialization

Use the command in [Table 4.50](#) to manage a drive initialization on the selected controller(s).

**Table 4.50 Drive Initialization**

<b>Convention</b>	<code>MegaCli -PDClear -Start  -Stop -ShowProg  -ProgDsply -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL</code>
<b>Description</b>	Manages initialization or displays initialization progress on a single controller, multiple controllers, or all controllers: -Start: Starts initialization on the selected drive(s). -Stop: Stops an ongoing initialization on the selected drive(s). -ShowProg: Displays the current progress percentage and time remaining for the initialization. This option is useful for running the application through scripts. -ProgDsply: Displays the ongoing clear progress. The routine continues to display the initialization progress until at least one initialization is completed or a key is pressed.

## 4.11.7 Rebuild a Drive

Use the command in [Table 4.51](#) to start or stop a rebuild on a drive and display the rebuild progress. When a drive in a RAID drive group fails, you can rebuild the drive by recreating the data that was stored on the drive before it failed.

**Table 4.51 Rebuild a Drive**

<b>Convention</b>	MegaCli -PDRbld -Start   -Stop   -ShowProg   -ProgDsply -PhysDrv [E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Manages a drive rebuild or displays the rebuild progress on a single controller, multiple controllers, or all controllers. Note that the drive must meet the capacity requirements before it can be rebuilt, and it must be part of a drive group: -Start: Starts a rebuild on the selected drive(s) and displays the rebuild progress (optional). -Stop: Stops an ongoing rebuild on the selected drive(s). -ShowProg: Displays the current progress percentage and time remaining for the rebuild. This option is useful for running the application through scripts. -ProgDsply: Displays the ongoing rebuild progress. This routine displays the rebuild progress until at least one initialization is completed or a key is pressed.

## 4.11.8 Locate the Drive(s) and Activate LED

Use the command in [Table 4.52](#) to locate the drive(s) for the selected controller(s) and activate the drive activity LED.

**Table 4.52 Locate Drive and Activate LED**

<b>Convention</b>	MegaCli -PDLocate -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Locates the drive(s) for the selected controller(s) and activates the drive activity LED.

## 4.11.9 Mark the Configured Drive as Missing

Use the command in [Table 4.53](#) to mark the configured drive as missing for the selected controller(s).

**Table 4.53 Mark Configured Drive as Missing**

<b>Convention</b>	MegaCli -PDMarkMissing -PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Marks the configured drive as missing for the selected controller(s).

## 4.11.10 Display the Drives in Missing Status

Use the command in [Table 4.53](#) to mark the configured drive as missing for the selected controller(s)

**Table 4.54 Display Drives in Missing Status**

<b>Convention</b>	MegaCli -PDGetMissing -aN -a0,1,2 -aALL												
<b>Description</b>	Displays the drive(s) in missing status. The format is:  <table><thead><tr><th>No</th><th>Row</th><th>Column</th><th>SizeExpected(MB)</th></tr></thead><tbody><tr><td>0</td><td>x</td><td>y</td><td>zzzzzzzz</td></tr><tr><td>...</td><td></td><td></td><td></td></tr></tbody></table> Where <i>x</i> is the index to the drive groups, <i>y</i> is the index to the drive in that drive group, and <i>zzzzzz</i> is the minimum capacity of the drive that can be used as a replacement.	No	Row	Column	SizeExpected(MB)	0	x	y	zzzzzzzz	...			
No	Row	Column	SizeExpected(MB)										
0	x	y	zzzzzzzz										
...													

## 4.11.11 Replace the Configured Drives and Start an Automatic Rebuild

Use the command in [Table 4.55](#) to replace configured drive(s) and start an automatic rebuild of the drive for the selected controller(s).

**Table 4.55 Replace Configured Drive(s) and Start Automatic Rebuild**

<b>Convention</b>	MegaCli -PDReplaceMissing -PhysDrv[E0:S0,E1:S1....] -ArrayX -RowY -aN
<b>Description</b>	Replaces the configured drives that are identified as missing and then starts an automatic rebuild.

## 4.11.12 Prepare the Unconfigured Drive for Removal

Use the command in [Table 4.56](#) to prepare the unconfigured drive(s) for removal from the selected controller(s).

**Table 4.56 Prepare Unconfigured Drive(s) for Removal**

<b>Convention</b>	MegaCli -PDPrpRmv [-Undo] - PhysDrv[E0:S0,E1:S1....] -aN -a0,1,2 -aALL
<b>Description</b>	Prepares unconfigured drive(s) for removal. The firmware will spin down this drive. The drive state is set to <i>unaffiliated</i> , which marks it as offline even though it is not a part of configuration. The <i>-Undo</i> option undoes this operation. If you select undo, the firmware marks this drive as <i>unconfigured good</i> .

### 4.11.13 Display Total Number of Drives

Use the command in [Table 4.57](#) to display the total number of drives attached to an controller. Drives can be attached directly or through enclosures.

**Table 4.57 Display Number of Drives Attached to an Controller**

<b>Convention</b>	MegaCli -PDGetNum -aN -a0,1,2 -aALL
<b>Description</b>	Displays the total number of drives attached to an controller. Drives can be attached directly or through enclosures. The return value is the number of drives.

### 4.11.14 Display List of Physical Devices

Use the command in [Table 4.58](#) to display a list of the physical devices connected to the selected controller(s).

**Table 4.58 Display List of Physical Devices Attached to Controller(s)**

<b>Convention</b>	MegaCli -PDLlist -aN -a0,1.. -aAll
<b>Description</b>	Displays information about all drives and other devices connected to the selected controller(s). This includes information such as the drive type, capacity (if a drive), serial number, and firmware version of the device. For SAS devices, this includes additional information such as the SAS address of the device. For SAS expanders, this includes additional information such as the number of drives connected to the expander.

### 4.11.15 Download Firmware to the Physical Devices

Use the command in [Table 4.59](#) to download firmware to the physical devices connected to the selected controller(s).

**Table 4.59 Download Firmware to the Physical Devices**

<b>Convention</b>	MegaCli -PdFwDownload - -PhysDrv[E0:S0,E1:S1....] --f <filename> -aN -a0,1,2 -aAll
<b>Description</b>	Flashes the firmware with the file specified at the command line. Firmware files used to flash the physical drive can be of any format. The CLI utility assumes that you provide a valid firmware image and flashes the same. The physical device has to do error checking. Firmware files in .dlp format can be flashed with the DOS version of the command tool only.

---

## 4.12 Enclosure-Related Options

The commands in this section are used for enclosures.

Use the command in [Table 4.60](#) to display enclosure information for selected controller(s).

**Table 4.60 Display Enclosure Information**

---

<b>Convention</b>	MegaCli -EncInfo -aN -a0,1,2 -aALL
<b>Description</b>	Displays information about the enclosure for the selected controller(s).

---

## 4.13 Flashing the Firmware

The options in this section describe the functionality of the existing flash application. The firmware flash options do not require input from the user.

### 4.13.1 Flash the Firmware with the ROM File

Use the command in [Table 4.61](#) to flash the firmware with the ROM file specified at the command line for the selected controller(s).

**Table 4.61 Flash Firmware with ROM File**

---

<b>Convention</b>	MegaCli -AdpFwFlash -f filename [-NoSigChk] [-NoVerChk]-aN -a0,1,2 -aALL
<b>Description</b>	<p>Flashes the firmware with the ROM file specified at the command line. The <code>-NoSigChk</code> option forces the application to flash the firmware even if the check word on the file does not match the required check word for the controller. This option flashes the firmware only if the existing firmware version on the controller is lower than the version on the ROM image.</p> <p>If you specify <code>-NoVerChk</code>, also, the application flashes the controller firmware without checking the version of the firmware image. The version check applies only to the firmware (APP.ROM) version.</p> <p>This command also supports the “Mode 0” flash functionality. For Mode 0 flash, the controller number is not valid. There are two possible methods:</p> <ul style="list-style-type: none"><li>• Select which controller to flash after the controllers are detected.</li><li>• Flash the firmware on all present controllers.</li></ul> <p>XML output data is generated by this option.</p>

---

## 4.13.2 Flash the Firmware in Mode 0 with the ROM File

Use the command in [Table 4.62](#) to flash the firmware in Mode 0 with the ROM file specified at the command line for the selected controller(s). This option is for DOS only.

**Table 4.62 Flash Firmware in Mode 0 with ROM File**

<b>Convention</b>	<code>MegaCli -AdpM0Flash -f filename</code>
<b>Description</b>	Flashes the firmware in Mode 0 with the ROM file listed on the command line. This option supports the Mode 0 flash functionality. For Mode 0 flash, the controller number is not valid. The method to handle this is to flash the firmware on all present controllers which are compatible with the image.

## 4.14 SAS Topology

The commands in this section are used to display SAS topology.

Use the command in [Table 4.63](#) to display the PHY connection information for physical PHY M on the selected controller(s). Each PHY can form one side of the physical link in a connection with a PHY on a different device. The physical link contains four wires that form two differential signal pairs. One differential pair transmits signals, and the other differential pair receives signals. Both differential pairs operate simultaneously and allow concurrent data transmission in both the receive and the transmit directions. PHYs are contained within ports.

A port can contain a single PHY or can contain multiple PHYs. A narrow port contains a single PHY, and a wide port contains multiple PHYs.

**Table 4.63 Display PHY Connection Information**

<b>Convention</b>	<code>MegaCli -PHYInfo -phyM -aN -a0,1,2 -aALL</code>
<b>Description</b>	Displays PHY connection information for physical PHY M on the controller(s).

---

## 4.15 Diagnostic-Related Options

The commands in this section are used to run diagnostic tests.

### 4.15.1 Start Controller Diagnostics

Use the command in [Table 4.64](#) to start the controller diagnostic for a set amount of time.

**Table 4.64 Start Diagnostics Setting**

<b>Convention</b>	<code>MegaCli -AdpDiag [val] -aN -a0,1,2 -aALL</code>
<b>Description</b>	Sets the amount of time for the controller diagnostic to run.

### 4.15.2 Start Battery Test

Use the command in [Table 4.65](#) to start the battery test. This command requires a system reboot.

**Table 4.65 Start Battery Test**

<b>Convention</b>	<code>MegaCli -AdpBatTest -aN -a0,1,2 -aALL</code>
<b>Description</b>	Starts the battery test. This command requires that you turn off the power to the system, and then turn on the power and reboot the system.

### 4.15.3 Start NVRAM Diagnostic

Use the command in [Table 4.66](#) to start the controller NVRAM diagnostic for a set amount of time. This option is for DOS only.

**Table 4.66 Start NVRAM Diagnostic**

<b>Convention</b>	<code>MegaCli -AdpNVRAM {-Read -Write -filename} -Clear [-StartOffset 0xXXXX] [-EndOffset 0xXXXX] aN</code>
<b>Description</b>	Starts the NVRAM diagnostic. -Read: Reads the content in NVRAM and writes the data to file <i>filename</i> . -Write: Reads data from file <i>filename</i> and writes to NVRAM. -Clear: Writes 0 to NVRAM at the specified range from start offset to end offset. -StartOffset/-EndOffset: Specifies the start offset and/or end offset in NVRAM. If you do not use the -StartOffset and -EndOffset options, the default StartOffset is 0 and the default EndOffset is the end of actual NVRAM size.

---

## 4.16 Miscellaneous Options

The commands in this section are used to display various information.

### 4.16.1 Display the MegaCLI Version

Use the command in [Table 4.67](#) to display the version number of the MegaCLI utility.

**Table 4.67 Display MegaCLI Version**

<b>Convention</b>	<code>MegaCli -v</code>
<b>Description</b>	Displays the version number of the MegaCLI utility.

### 4.16.2 Display Help for MegaCLI

Use the command in [Table 4.68](#) to display help information for the MegaCLI utility.

**Table 4.68 Display Help for MegaCLI**

<b>Convention</b>	<code>MegaCli -h -Help ?</code>
<b>Description</b>	Displays help for the MegaCLI utility.

# Appendix

## Glossary

---

<b>access policy</b>	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are <i>Read/Write</i> , <i>Read Only</i> , or <i>Blocked</i> .
<b>alarm enabled</b>	A controller property that indicates whether the controller's onboard alarm is enabled.
<b>alarm present</b>	A controller property that indicates whether the controller has an onboard alarm. If present and enabled, the alarm is sounded for certain error conditions.
<b>array</b>	See <i>drive group</i> .
<b>BBU present</b>	A controller property that indicates whether the controller has an onboard battery backup unit to provide power in case of a power failure.
<b>BGI rate</b>	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
<b>BIOS</b>	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.
<b>cache</b>	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.

<b>cache flush interval</b>	A controller property that indicates how often the data cache is flushed.
<b>caching</b>	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache may temporarily store data in accordance with its write back policies.
<b>capacity</b>	A property that indicates the amount of storage space on a drive or virtual drive.
<b>coerced capacity</b>	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4 Gbyte drive from one manufacturer may be 4,196 Mbytes, and a 4 Gbyte from another manufacturer may be 4,128 Mbytes. These drives could be coerced to a usable capacity of 4,088 Mbytes each for use in a drive group in a storage configuration.
<b>coercion mode</b>	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.
<b>consistency check</b>	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
<b>consistency check rate</b>	The rate at which consistency check operations are run on a computer system.
<b>controller</b>	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection. MegaRAID Storage Manager software runs on LSI SAS controllers.

<b>copyback</b>	<p>The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). The copyback operation can be run automatically or manually.</p> <p>Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.</p>
<b>current write policy</b>	<p>A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode.</p> <ul style="list-style-type: none"> <li>• In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.</li> <li>• In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.</li> </ul>
<b>default write policy</b>	<p>A virtual drive property indicating whether the default write policy is Write Through or Write Back. In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.</p>
<b>device ID</b>	<p>A controller or drive property indicating the manufacturer-assigned device ID.</p>
<b>device port count</b>	<p>A controller property indicating the number of ports on the controller.</p>
<b>drive cache policy</b>	<p>A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.</p>
<b>drive group</b>	<p>A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.</p>

<b>drive state</b>	<p>A drive property indicating the status of the drive. A drive can be in one of the following states:</p> <ul style="list-style-type: none"> <li>• Unconfigured Good: A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare.</li> <li>• Hot Spare: A drive that is configured as a hot spare.</li> <li>• Online: A drive that can be accessed by the RAID controller and will be part of the virtual drive.</li> <li>• Rebuild: A drive to which data is being written to restore full redundancy for a virtual drive.</li> <li>• Failed: A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.</li> <li>• Unconfigured Bad: A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.</li> <li>• Missing: A drive that was Online, but which has been removed from its location.</li> <li>• Offline: A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned.</li> <li>• None: A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.</li> </ul>
<b>drive subsystem</b>	<p>A collection of drives and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the drives can attach directly to a system I/O bus controller.</p>
<b>drive type</b>	<p>A drive property indicating the characteristics of the drive.</p>
<b>fast initialization</b>	<p>A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.</p>
<b>fault tolerance</b>	<p>The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. LSI SAS RAID controllers provides fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.</p>

<b>firmware</b>	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 drive or from a network and then passes control to the operating system.
<b>foreign configuration</b>	A RAID configuration that already exists on a replacement set of drives that you install in a computer system. MegaRAID Storage Manager software allows you to import the existing configuration to the RAID controller, or you can clear the configuration so you can create a new one.
<b>formatting</b>	The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.
<b>hole</b>	In MegaRAID Storage Manager, a <i>hole</i> is a block of empty space in a drive group that can be used to define a virtual drive.
<b>host interface</b>	A controller property indicating the type of interface used by the computer host system: for example, <i>PCIX</i> .
<b>host port count</b>	A controller property indicating the number of host data ports currently in use.
<b>host system</b>	Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.
<b>hot spare</b>	<p>A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller.</p> <p>When a drive fails, MegaRAID Storage Manager software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.</p>

<b>initialization</b>	The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.
<b>IO policy</b>	A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)
<b>load-balancing</b>	A method of spreading work between two or more computers, network links, CPUs, drives, or other resources. Load balancing is used to maximize resource use, throughput, or response time.
<b>media error count</b>	A drive property indicating the number of errors that have been detected on the drive media.
<b>migration</b>	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.
<b>mirroring</b>	The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.
<b>multipathing</b>	The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.
<b>name</b>	A virtual drive property indicating the user-assigned name of the virtual drive.

<b>non-redundant configuration</b>	A RAID 0 virtual drive with data striped across two or more drives but without drive mirroring or parity. This provides for high data throughput but offers no protection in case of a drive failure.
<b>NVRAM</b>	Acronym for non-volatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.
<b>NVRAM present</b>	A controller property indicating whether an NVRAM is present on the controller.
<b>NVRAM size</b>	A controller property indicating the capacity of the controller's NVRAM.
<b>offline</b>	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.
<b>patrol read</b>	A process that checks the drives in a storage configuration for drive errors that could lead to drive failure and lost data. The patrol read operation can find and sometimes fix any potential problem with drives prior to host access. This enhances overall system performance because error recovery during a normal I/O operation may not be necessary.
<b>patrol read rate</b>	The user-defined rate at which patrol read operations are run on a computer system.
<b>product info</b>	A drive property indicating the vendor-assigned model number of the drive.
<b>product name</b>	A controller property indicating the manufacturing name of the controller.
<b>RAID</b>	A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data. A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.
<b>RAID 0</b>	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.

<b>RAID 00</b>	Uses data striping on two or more drives in a spanned drive group to provide high data throughput, especially for large files in an environment that requires no data redundancy.
<b>RAID 1</b>	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
<b>RAID 5</b>	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
<b>RAID 6</b>	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
<b>RAID 10</b>	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
<b>RAID 50</b>	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
<b>RAID 60</b>	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.
<b>RAID level</b>	A virtual drive property indicating the RAID level of the virtual drive. LSI SAS controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
<b>raw capacity</b>	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.

<b>read policy</b>	A controller attribute indicating the current Read Policy mode. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode, read ahead capability is disabled. In Adaptive Read Ahead mode, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to No Read Ahead mode.
<b>rebuild</b>	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
<b>rebuild rate</b>	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.
<b>reclaim virtual drive</b>	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click the <b>Reclaim</b> button, the individual drives are removed from the virtual drive configuration.
<b>reconstruction rate</b>	The user-defined rate at which a reconstruction operation is carried out.
<b>redundancy</b>	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
<b>redundant configuration</b>	A virtual drive that has redundant data on drives in the drive group that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive. A redundant configuration protects the data in case a drive fails in the configuration.
<b>revertible hot spare</b>	When you use the Replace Member procedure, after data is copied from a hot spare to a new drive, the hot spare reverts from a rebuild drive to its original hot spare status.
<b>revision level</b>	A drive property that indicates the revision level of the drive's firmware.

<b>SAS</b>	Acronym for Serial Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
<b>SATA</b>	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
<b>SCSI device type</b>	A drive property indicating the type of the device, such as drive.
<b>serial no.</b>	A controller property indicating the manufacturer-assigned serial number.
<b>strip size</b>	The portion of a stripe that resides on a single drive in the drive group.
<b>stripe size</b>	A virtual drive property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. The user can select the stripe size.
<b>striping</b>	A technique used to write data across all drives in a virtual drive. Each stripe consists of consecutive virtual drive data addresses that are mapped in fixed-size units to each drive in the virtual drive using a sequential pattern. For example, if the virtual drive includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
<b>subvendor ID</b>	A controller property that lists additional vendor ID information about the controller.
<b>uncorrectable error count</b>	A controller property that lists the number of uncorrectable errors detected on drives connected to the controller. If the error count reaches a certain level, a drive will be marked as failed.
<b>vendor ID</b>	A controller property indicating the vendor-assigned ID number of the controller.

<b>vendor info</b>	A drive property listing the name of the vendor of the drive.
<b>virtual drive</b>	A storage unit created by a RAID controller from one or more drives. Although a virtual drive may be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive may retain redundant data in case of a drive failure.
<b>virtual drive state</b>	A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.
<b>write-back</b>	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive 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, and elapsed time from the last cache flush.
<b>write policy</b>	See <i>Default Write Policy</i> .
<b>write-through</b>	In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.