



MegaRAID®
SAS Software
User Guide

51530-00
Rev. B
July 2011



Revision History

Version and Date	Description of Changes
51530-00, Rev B, July 2011	<ul style="list-style-type: none">• Updated the guide with VMware 5.0 information.• Updated the guide with CacheCade Pro 2.0 SSD Read/Write Caching software content.
51530-00A, Rev A, May 2011	<ul style="list-style-type: none">• Updated screen shots wherever required in the document.• Added around 53 events from 0x0189 to 0x01bd in Appendix A.• Updated the content in the Section 5.15.15, Download Firmware to the Physical Devices.• Modified content in Chapter 4, WebBIOS Configuration Utility.• Updated content in Chapter 5, MegaRAID Command Tool.• Updated content in Chapter 6, MegaRAID Storage Manager Overview and Installation.• Updated content in Chapter 7, MegaRAID Storage Manager Window and Menus.• Updated content in Chapter 8, Configuration.• Updated content in Chapter 9, Monitoring Controllers and Its Attached Devices.• Updated content in Chapter 10, Maintaining and Managing Storage Configurations.• Updated content in Chapter 11, Using MegaRAID Advanced Software.• Updated content in Appendix A, Events and Messages.• Created a new appendix - Appendix D, History of Technical Changes.• Updated the name of the CacheCade SSD Caching software.

NOTE: For a history of all technical changes made to this guide for the previous releases, refer to [Appendix D, History of Technical Changes](#).

LSI and the LSI logo are trademarks or registered trademarks of LSI Corporation or its subsidiaries. All other brand and product names may be trademarks of their respective companies.

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

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 LSI Corporation.

Corporate Headquarters
Milpitas, CA
800-372-2447

Email
globalsupport@lsi.com

Website
www.lsi.com

Document Number: 51530-00, Rev. B
Copyright © 2011 LSI Corporation
All Rights Reserved

Table of Contents

Chapter 1: Overview	13
1.1 SAS Technology	13
1.2 Serial-Attached SCSI Device Interface	14
1.3 Serial ATA III Features	15
1.4 Solid State Drive Features	15
1.4.1 SSD Guard	16
1.5 Dimmer Switch Features	16
1.6 UEFI 2.0 Support	16
1.7 Configuration Scenarios	16
1.7.1 Valid Drive Mix Configurations with HDDs and SSDs	18
1.8 Technical Support	19
Chapter 2: Introduction to RAID	21
2.1 RAID Description	21
2.2 RAID Benefits	21
2.3 RAID Functions	21
2.4 Components and Features	22
2.4.1 Drive Group	22
2.4.2 Virtual Drive	22
2.4.3 Fault Tolerance	22
2.4.4 Consistency Check	23
2.4.5 Copyback	23
2.4.6 Background Initialization	24
2.4.7 Patrol Read	25
2.4.8 Disk Striping	25
2.4.9 Disk Mirroring	26
2.4.10 Parity	26
2.4.11 Disk Spanning	27
2.4.12 Hot Spares	28
2.4.13 Disk Rebuilds	29
2.4.14 Rebuild Rate	30
2.4.15 Hot Swap	30
2.4.16 Drive States	30
2.4.17 Virtual Drive States	31
2.4.18 Beep Codes	31
2.4.19 Enclosure Management	32
2.5 RAID Levels	32
2.5.1 Summary of RAID Levels	32
2.5.2 Selecting a RAID Level	33
2.5.3 RAID 0	33
2.5.4 RAID 1	34
2.5.5 RAID 5	34
2.5.6 RAID 6	35
2.5.7 RAID 00	36
2.5.8 RAID 10	37
2.5.9 RAID 50	38
2.5.10 RAID 60	39

2.6 RAID Configuration Strategies	40
2.6.1 Maximizing Fault Tolerance	41
2.6.2 Maximizing Performance	42
2.6.3 Maximizing Storage Capacity	43
2.7 RAID Availability	44
2.7.1 RAID Availability Concept	44
2.8 Configuration Planning	45
2.9 Number of Drives	45
2.9.1 Drive Group Purpose	45
Chapter 3: SafeStore Disk Encryption	47
3.1 Overview	47
3.2 Purpose and Benefits	47
3.3 Terminology	48
3.4 Workflow	48
3.4.1 Enable Security	48
3.4.2 Change Security	49
3.4.3 Create Secure Virtual Drives	50
3.4.4 Import a Foreign Configuration	50
3.5 Instant Secure Erase	50
Chapter 4: WebBIOS Configuration Utility	53
4.1 Overview	53
4.2 Starting the WebBIOS configuration utility	53
4.3 WebBIOS configuration utility Main Dialog Options	54
4.4 Managing Software Licensing	56
4.4.1 Managing MegaRAID Advanced Software Options	56
4.4.2 Reusing the Activation Key	60
4.4.3 Managing Advanced Software Summary	60
4.4.4 Activating an Unlimited Key Over a Trial Key	61
4.4.5 Activating a Trial Software	61
4.4.6 Activating an Unlimited Key	62
4.4.7 Securing MegaRAID Advanced Software	63
4.4.8 Confirm Re-hosting Process	63
4.4.9 Re-hosting Process Complete	65
4.5 Creating a Storage Configuration	66
4.5.1 Using Automatic Configuration	69
4.5.2 Using Manual Configuration	69
4.6 CacheCade Configuration	111
4.6.1 Creating a CacheCade 2.0 SSD Read Caching Configuration	111
4.6.2 Creating a CacheCade Pro 2.0 SSD Read/Write Caching Configuration	117
4.7 Selecting SafeStore Encryption Services Security Options	126
4.7.1 Enabling the Security Key Identifier, Security Key, and Password	126
4.7.2 Enabling Drive Security using EKM	132
4.7.3 Changing the Security Key Identifier, Security Key, and Pass Phrase	133
4.7.4 Change Security from EKM to LKM	134
4.7.5 Changing Security from LKM to EKM	139
4.7.6 Disabling the Drive Security Settings	140
4.8 Viewing and Changing Device Properties	142
4.8.1 Viewing Controller Properties	142
4.8.2 Viewing Virtual Drive Properties, Policies, and Operations	149
4.8.3 Viewing Drive Properties	150

4.8.4 Shield State	152
4.8.5 Viewing and Changing Battery Backup Unit Information	155
4.8.6 Managing Link Speed	160
4.8.7 Viewing Enclosure Properties	162
4.8.8 SSD Disk Cache Policy	167
4.8.9 Emergency Hotspare	167
4.8.10 Emergency Hotspare for Controllers	168
4.9 Viewing and Expanding a Virtual Drive	172
4.10 Suspending and Resuming Virtual Drive Operations	173
4.11 Using MegaRAID Recovery	175
4.11.1 Recovery Scenarios	175
4.11.2 Enabling the Recovery Advanced Software	176
4.11.3 Creating Snapshots and Views	178
4.11.4 Creating Concurrent Snapshots	182
4.11.5 Selecting the Snapshot Settings	184
4.11.6 Viewing Snapshot Properties	186
4.11.7 Restoring a Virtual Drive by Rolling Back to a Snapshot	187
4.11.8 Cleaning Up a Snapshot Repository	189
4.12 Non-SED Secure Erase	191
4.12.1 Erasing a Non - SED Physical Drive	191
4.12.2 Virtual Drive Erase	195
4.13 Viewing System Event Information	198
4.14 Managing Configurations	199
4.14.1 Running a Consistency Check	199
4.14.2 Deleting a Virtual Drive	199
4.14.3 Importing or Clearing a Foreign Configuration	200
4.14.4 Importing Foreign Configurations	204
4.14.5 Import Foreign Drives in EKM/EKM Secured Locked Drives	204
4.14.6 Importing Foreign Drives for LKM-Secured Locked Drives	205
4.14.7 Importing Foreign Drives in LKM Mode EKM-Secured Locked Drives	205
4.14.8 Migrating the RAID Level of a Virtual Drive	206
4.14.9 New Drives Attached to a MegaRAID Controller	208
4.15 WebBIOS Dimmer Switch	208
4.15.1 Power-Save Mode	214
4.15.2 Power Save Settings-Advanced	215
4.15.3 Power-Save While Creating Virtual Drives	216
Chapter 5: MegaRAID Command Tool	217
5.1 Product Overview	217
5.2 Novell NetWare, SCO, Solaris, FreeBSD, and MS-DOS Operating System Support	218
5.3 Command Line Abbreviations and Conventions	219
5.3.1 Abbreviations Used in the Command Line	219
5.3.2 Conventions	219
5.4 Pre-boot MegaCLI	220
5.5 CacheCade Related Options	221
5.5.1 Create a Solid State Drive Cache Drive to Use as Secondary Cache	222
5.5.2 Delete a Solid State Drive Cache Drive	222
5.5.3 Associate/Disassociate Virtual Drives	222
5.5.4 Display CacheCade Pro 2.0 Configurations on a Controller	222
5.5.5 Create a RAID Drive Group for CacheCade Pro 2.0 from All Unconfigured Good Drives	223

5.5.6 Remove Blocked Access on a Virtual Drive	224
5.5.7 Create RAID 0 Configuration with SSD Caching	224
5.5.8 Create a RAID Level 10, 50, 60 (spanned) Configuration with SSD Caching	225
5.5.9 Delete Virtual Drives with SSD Caching	226
5.5.10 Clear Configurations on CacheCade Pro 2.0 Virtual Drives	226
5.5.11 Create a CacheCade Pro 2.0 Virtual Drive with RAID level and Write Policy	226
5.6 Software License Key	226
5.7 SafeStore Security Options	227
5.7.1 Use Instant Secure Erase on a Physical Drive	228
5.7.2 Secure Data on a Virtual Drive	228
5.7.3 Destroy the Security Key	228
5.7.4 Create a Security Key	229
5.7.5 Create a Drive Security Key	229
5.7.6 Change the Security Key	230
5.7.7 Get the Security Key ID	230
5.7.8 Set the Security Key ID	230
5.7.9 Verify the Security Key	231
5.8 Controller Property-Related Options	231
5.8.1 Display Controller Properties	231
5.8.2 Display Number of Controllers Supported	231
5.8.3 Enable or Disable Automatic Rebuild	231
5.8.4 Flush Controller Cache	232
5.8.5 Set Controller Properties	232
5.8.6 Display Specified Controller Properties	234
5.8.7 Set Factory Defaults	234
5.8.8 Set SAS Address	234
5.8.9 Set Time and Date on Controller	235
5.8.10 Display Time and Date on Controller	235
5.8.11 Get Connector Mode	235
5.8.12 Set Connector Mode	235
5.9 Patrol Read-Related Controller Properties	236
5.9.1 Set Patrol Read Options	236
5.9.2 Set Patrol Read Delay Interval	237
5.10 BIOS-Related Properties	237
5.10.1 Set or Display Bootable Virtual Drive ID	237
5.10.2 Select BIOS Status Options	237
5.11 Battery Backup Unit-Related Properties	238
5.11.1 Display BBU Information	238
5.11.2 Display BBU Status Information	239
5.11.3 Display BBU Capacity	240
5.11.4 Display BBU Design Parameters	240
5.11.5 Display Current BBU Properties	241
5.11.6 Start BBU Learning Cycle	241
5.11.7 Place Battery in Low-Power Storage Mode	241
5.11.8 Set BBU Properties	242
5.12 Options for Displaying Logs Kept at the Firmware Level	242
5.12.1 Event Log Management	242
5.12.2 Set BBU Terminal Logging	243
5.13 Configuration-Related Options	243
5.13.1 Create a RAID Drive Group from All Unconfigured-Good Drives	243
5.13.2 Add RAID 0, 1, 5, or 6 Configuration	245
5.13.3 Add RAID 10, 50, or 60 Configuration	246
5.13.4 Clear the Existing Configuration	246

5.13.5 Save the Configuration on the Controller	246
5.13.6 Restore the Configuration Data from File	247
5.13.7 Manage Foreign Configuration Information	247
5.13.8 Delete Specified Virtual Drives	247
5.13.9 Display the Free Space	248
5.14 Virtual Drive-Related Options	248
5.14.1 Display Virtual Drive Information	248
5.14.2 Change the Virtual Drive Cache and Access Parameters	248
5.14.3 Display the Virtual Drive Cache and Access Parameters	249
5.14.4 Manage Virtual Drives Initialization	249
5.14.5 Manage a Consistency Check	249
5.14.6 Schedule a Consistency Check	250
5.14.7 Manage a Background Initialization	250
5.14.8 Perform a Virtual Drive Reconstruction	251
5.14.9 Display Information about Virtual Drives and Drives	252
5.14.10 Display the Bad Block Table	252
5.14.11 Display the Number of Virtual Drives	252
5.14.12 Clear the LDBBM Table Entries	253
5.14.13 Display the List of Virtual Drives with Preserved Cache	253
5.14.14 Discard the Preserved Cache of a Virtual Drives	253
5.14.15 Expand a Virtual Drive	253
5.15 Drive-Related Options	254
5.15.1 Display Drive Information	254
5.15.2 Set the Drive State to Online	254
5.15.3 Set the Drive State to Offline	254
5.15.4 Change the Drive State to Unconfigured-Good	255
5.15.5 Change the Drive State	255
5.15.6 Manage a Drive Initialization	256
5.15.7 Rebuild a Drive	256
5.15.8 Locate the Drives and Activate LED	257
5.15.9 Mark the Configured Drive as Missing	257
5.15.10 Display the Drives in Missing Status	257
5.15.11 Replace the Configured Drives and Start an Automatic Rebuild	258
5.15.12 Prepare the Unconfigured Drive for Removal	258
5.15.13 Display Total Number of Drives	258
5.15.14 Display List of Physical Devices	259
5.15.15 Download Firmware to the Physical Devices	259
5.15.16 Configure All Free Drives into a RAID 0, 1, 5, or 6 Configuration for a Specific Controller	260
5.15.17 Set the Mapping Mode of the Drives to the Selected Controllers	261
5.15.18 Secure Erase for Virtual Drives and Physical Drives	261
5.15.19 Perform the Copyback Operation on the Selected Drive	261
5.16 Enclosure-Related Options	262
5.16.1 Display Enclosure Information	262
5.16.2 Display Enclosure Status	263
5.16.3 Upgrading the Firmware without Restarting	263
5.17 Flashing the Firmware	263
5.17.1 Flash the Firmware with the ROM File	263
5.17.2 Flash the Firmware in Mode 0 with the ROM File	264
5.18 SAS Topology	264
5.19 Diagnostic-Related Options	264
5.19.1 Start Controller Diagnostics	264
5.19.2 Perform a Full Stroke Seek Test	265

5.19.3 Start Battery Test	265
5.20 Recovery (Snapshot)-Related Options	265
5.20.1 Enable the Snapshot Feature	265
5.20.2 Disable the Snapshot Feature	266
5.20.3 Take a Snapshot of a Volume	266
5.20.4 Set the Snapshot Properties	266
5.20.5 Delete a Snapshot	267
5.20.6 Create a View	267
5.20.7 Delete a View	267
5.20.8 Roll back to an Older Snapshot	268
5.20.9 Display Snapshot and View Information	268
5.20.10 Clean the Recoverable Free Space on the Drives in a Virtual Drive	268
5.20.11 Display the Information for a Specific View	269
5.20.12 Enable the Snapshot Scheduler	269
5.21 FastPath-Related Options	269
5.22 Dimmer Switch-Related Options	270
5.22.1 Display Selected Adapter Properties	270
5.22.2 Sets the Properties on the Selected Adapter	271
5.22.3 Displays the Power-Saving Level on the Virtual Disk	271
5.22.4 Displays about Adding a RAID Level to a Specified Adapter	272
5.22.5 Create a RAID Level	273
5.22.6 Add the Unconfigured Drive to a Specified Adapter	274
5.22.7 Displays the Cache and Access Policies	275
5.23 Performance Monitoring Options	275
5.23.1 Starting Performance Data Collection	275
5.23.2 Stopping Performance Data Collection	276
5.23.3 Saving Performance Data	276
5.24 Miscellaneous Commands	276
5.24.1 Display the Version	276
5.24.2 Display the MegaCLI Version	276
5.24.3 Display Help for MegaCLI	277
5.24.4 Display Summary Information	277
Chapter 6: MegaRAID Storage Manager Overview and Installation	279
6.1 Overview	279
6.1.1 Creating Storage Configurations	279
6.1.2 Monitoring Storage Devices	279
6.1.3 Maintaining Storage Configurations	279
6.2 Hardware and Software Requirements	280
6.3 Prerequisites to Running MegaRAID Storage Manager Remote Administration	280
6.4 Installing MegaRAID Storage Manager	281
6.4.1 Prerequisite for MegaRAID Storage Manager Installation	281
6.4.2 Installing MegaRAID Storage Manager Software on Microsoft Windows	281
6.4.3 Installing MegaRAID Storage Manager for the Solaris SPARC Operating System	286
6.4.4 Uninstalling MegaRAID Storage Manager Software for Solaris SPARC Operating System	286
6.4.5 Installing MegaRAID Storage Manager for Linux Operating System	287
6.4.6 Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.0 x64 Operating System	288
6.4.7 Linux Error Messages	288
6.4.8 Kernel Upgrade	289
6.4.9 Uninstalling MegaRAID Storage Manager Software on Linux	289
6.4.10 MegaRAID Storage Manager Customization	290
6.5 MegaRAID Storage Manager Support and Installation on VMware	291
6.5.1 Pre-requisites for Installing MegaRAID Storage Manager for VMware	291

6.5.2 Installing MegaRAID Storage Manager on VMware ESX (VMware Classic)	291
6.5.3 Uninstalling MegaRAID Storage Manager for VMware	291
6.5.4 MegaRAID Storage Manager Support on the VMware ESXi Operating System	292
6.5.5 Limitations	292
6.6 Installing and Configuring a CIM Provider	294
6.6.1 Installing a CIM SAS Storage Provider on the Linux Operating System	294
6.6.2 Installing a CIM SAS Storage Provider on Windows	295
6.7 Installing and Configuring an SNMP Agent	295
6.7.1 Prerequisite for LSI SNMP Agent RPM Installation	296
6.7.2 Prerequisite for Installing SNMP Agent on Linux Server	296
6.7.3 Installing and Configuring an SNMP Agent on a Linux Operating System	296
6.7.4 Installing and Configuring an SNMP Agent on the Solaris Operating System	298
6.7.5 Installing an SNMP Agent on the Windows Operating System	300
6.8 MegaRAID Storage Manager Support and Installation on the Solaris 10 Operating System	302
6.8.1 Installing MegaRAID Storage Manager Software for the Solaris 10 Operating System	302
6.8.2 Uninstalling MegaRAID Storage Manager Software for the Solaris 10 Operating System	302
6.9 Installing MegaCLI for VMware 5.0	302
6.10 MegaRAID Storage Manager Remotely Connecting to VMware ESX	303
6.11 Prerequisites to Running MegaRAID Storage Manager Remote Administration	303
Chapter 7: MegaRAID Storage Manager Window and Menus	305
7.1 Starting the MegaRAID Storage Manager Software	305
7.2 MegaRAID Storage Manager Main Menu	309
7.2.1 Dashboard, Physical View, Logical View	310
7.2.2 Shield State	314
7.2.3 Shield State Physical View	314
7.2.4 Logical View Shield State	315
7.2.5 Viewing the Physical Drive Properties	315
7.2.6 Viewing Server Profile of a Drive in Shield State	316
7.2.7 Displaying the Virtual Drive Properties	317
7.2.8 Emergency HotSpare Support	320
7.2.9 SSD Disk Cache Policy	322
7.2.10 Non-SED Secure Erase Support	324
7.2.11 Rebuild Write Cache	329
7.2.12 Background Suspend or Resume Support	330
7.2.13 Enclosure Properties	331
7.3 Monitoring Battery Backup Units	331
7.3.1 Properties and Graphical View Tabs	334
7.3.2 Event Log Panel	335
7.3.3 Menu Bar	335
Chapter 8: Configuration	337
8.1 Creating a New Storage Configuration	337
8.1.1 Selecting Virtual Drive Settings	337
8.1.2 Optimum Controller Settings for CacheCade	339
8.1.3 Optimum Controller Settings for FastPath	339
8.1.4 Creating a Virtual Drive Using Simple Configuration	339
8.1.5 Creating a Virtual Drive using Advanced Configuration	344
8.2 Converting JBOD Drives to Unconfigured Good	350
8.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Window	351

8.3 Adding Hot Spare Drives	352
8.4 Changing Adjustable Task Rates	353
8.5 Changing Power Settings	355
8.5.1 Enhanced Dimmer Switch Power Settings	356
8.5.2 Power Save Settings - Advanced	358
8.5.3 Automatically Spin Up Drives	359
8.5.4 Power-Save Mode	360
8.5.5 Power Save Mode - SSD Drives	361
8.6 Changing Virtual Drive Properties	361
8.7 Changing a Virtual Drive Configuration	363
8.7.1 Accessing the Modify Drive Group Wizard	363
8.7.2 Adding a Drive or Drives to a Configuration	365
8.7.3 Removing a Drive from a Configuration	367
8.7.4 Replacing a Drive	367
8.7.5 Migrating the RAID Level of a Virtual Drive	368
8.7.6 New Drives Attached to a MegaRAID Controller	370
8.8 Deleting a Virtual Drive	371
Chapter 9: Monitoring Controllers and Its Attached Devices	373
9.1 Alert Delivery Methods	373
9.1.1 Vivaldi Log / MegaRAID Storage Manager Log	374
9.1.2 System Log	375
9.1.3 Pop-up Notification	375
9.1.4 E-mail Notification	375
9.2 Configuring Alert Notifications	376
9.3 Editing Alert Delivery Methods	378
9.4 Changing Alert Delivery Methods for Individual Events	379
9.5 Changing the Severity Level for Individual Events	380
9.6 Rollback to Default Individual Event Configuration	381
9.7 Entering or Editing the Sender Email Address and SMTP Server	381
9.8 Authenticating the SMTP Server	382
9.9 Adding Email Addresses of Recipients of Alert Notifications	383
9.10 Testing Email Addresses of Recipients of Alert Notifications	383
9.11 Removing Email Addresses of Recipients of Alert Notifications	384
9.12 Saving Backup Configurations	384
9.13 Loading Backup Configurations	384
9.14 Monitoring Server Events	385
9.15 Monitoring Controllers	385
9.16 Monitoring Drives	386
9.17 Running a Patrol Read	387
9.18 Monitoring Virtual Drives	389
9.19 Monitoring Enclosures	390
9.19.1 Monitoring Battery Backup Units	391
9.20 Battery Learn Cycle	392
9.20.1 Setting Learn Cycle Properties	392
9.20.2 Starting a Learn Cycle Manually	393
9.21 Monitoring Rebuilds and Other Processes	394

Chapter 10: Maintaining and Managing Storage Configurations	397
10.1 Initializing a Virtual Drive	397
10.1.1 Running a Group Initialization	398
10.2 Running a Consistency Check	398
10.2.1 Setting the Consistency Check Settings	399
10.2.2 Scheduling a Consistency Check	400
10.2.3 Running a Group Consistency Check	402
10.3 Scanning for New Drives	403
10.4 Rebuilding a Drive	403
10.4.1 New Drives Attached to a MegaRAID Controller	404
10.5 Making a Drive Offline or Missing	405
10.6 Removing a Drive	405
10.7 Upgrading the Firmware	406
Chapter 11: Using MegaRAID Advanced Software	407
11.1 MegaRAID Advanced Software	407
11.2 Recovery Advanced Software	407
11.2.1 MegaRAID Software Licensing	407
11.2.2 Managing MegaRAID Advanced Software	408
11.2.3 Activation Key	410
11.2.4 Advanced MegaRAID Software Status Summary	411
11.2.5 Application Scenarios and Messages	412
11.2.6 Activating an Unlimited Key Over a Trial Key	413
11.2.7 Configuring Key Vault (Re-hosting process)	416
11.2.8 Re-hosting Complete	418
11.2.9 Deactivate Trial Software	419
11.2.10 MegaRAID Recovery	420
11.2.11 Recovery Scenarios	420
11.2.12 Enabling the Recovery Advanced Software	421
11.2.13 Snapshot Repository	422
11.2.14 Selecting the Virtual Drive	424
11.2.15 Scheduling Snapshots	425
11.2.16 Editing Snapshots	427
11.2.17 Snapshot Base Details	429
11.2.18 Manage Snapshots	430
11.2.19 Editing Schedule	431
11.2.20 Advanced Settings	432
11.2.21 Create View Using Manage Snapshots Wizard	433
11.2.22 Viewing Snapshot Details	434
11.2.23 No View Details for Snapshot	435
11.2.24 No Snapshot Schedule	436
11.2.25 Graphical Representation of Repository Virtual Drive	437
11.2.26 Deleting a Snapshot	437
11.3 Disabling MegaRAID Recovery	438
11.4 CacheCade Advanced Software	438
11.4.1 Using the CacheCade 2.0 SSD Read Caching Advanced Software	439
11.4.2 Using the CacheCade Pro 2.0 SSD Read/Write Caching Software	442
11.5 FastPath Advanced Software	451
11.5.1 Setting FastPath Options	451
11.6 LSI SafeStore Encryption Services	452

11.6.1 Enabling Drive Security using EKM	452
11.6.2 Supporting EKM Mode	453
11.6.3 Change Security Settings- LKM	455
11.6.4 Change Security Settings - EKM	457
11.6.5 Importing Foreign Drives	458
11.6.6 Importing Foreign Drives to LKM	459
11.6.7 Importing Foreign Drives to EKM	459
11.6.8 Importing Foreign Drives to EKM	460
11.6.9 Enabling Drive Security using LKM	461
11.6.10 Changing the Drive Security Settings	465
11.6.11 Disabling Drive Security	466
11.6.12 Importing or Clearing a Foreign Configuration	468
11.7 Managing Link Speed	471
Appendix A: Events and Messages	473
A.1 Error Levels	473
A.2 Event Messages	473
Appendix B: MegaCLI Error Messages	487
B.1 Error Messages and Descriptions	487
Appendix C: Glossary	491
Appendix D: History of Technical Changes	499

Chapter 1

Overview

This chapter provides an overview of this guide, which 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 describes how to use the MegaRAID Storage Manager™ software, the WebBIOS™ configuration utility, and the MegaRAID command line interface (CLI).

This chapter documents the SAS technology, Serial ATA (SATA) technology, MegaRAID CacheCade™ 2.0 SSD Read Caching software, SSD Guard™, Dimmer Switch™, UEFI 2.0, configuration scenarios, and drive types. Other features such as FastPath and SafeStore are described in other chapters of this guide.

NOTE: This guide does not include the latest CacheCade and Enterprise Key Management System (EKMS) features.

1.1 SAS Technology

The MegaRAID 6Gb/s SAS RAID controllers are high-performance intelligent PCI Express-to-SAS/Serial ATA II controllers with RAID control capabilities. MegaRAID 6Gb/s 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 6Gb/s 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 CacheCade 2.0 SSD Read Caching software devices within the same storage infrastructure. These devices bring individual characteristics that make each of these 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 combine SAS drives and SATA drives within the same *virtual drives*. Although you can mix drives, LSI strongly discourages this practice; this applies to both HDDs and CacheCade 2.0 SSD Read Caching software.

MegaRAID 6Gb/s 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 2.1*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, version 3.0*. 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 or SATA III devices using the following protocols:

- SAS Serial SCSI Protocol (SSP), which enables communication with other SAS devices
- SATA III, which enables communication with other SATA III 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 III 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 the 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 III Features

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

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

- Supports SATA III data transfers of 6Gb/s
- Supports STP data transfers of 6Gb/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 Drive Features

The MegaRAID firmware supports the use of SSDs as standard drives and/or additional controller cache, referred to as CacheCade 2.0 SSD Read Caching software. SSD drives are expected to behave like SATA or SAS HDDs except for the following:

- 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

NOTE: Support for SATA SSD drives applies only to those drives that support ATA-8 ACS compliance.

You can choose whether to allow a virtual drive to consist of both CacheCade 2.0 SSD Read Caching software devices and HDDs. For a virtual drive that consists of CacheCade 2.0 SSD Read Caching software only, you can choose whether to allow SAS CacheCade 2.0 SSD Read Caching software drives and SATA CacheCade 2.0 SSD Read Caching software drives in that virtual drive. For virtual drives that have both CacheCade 2.0 SSD Read Caching software and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA CacheCade 2.0 SSD Read Caching software devices in various combinations.

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

1.4.1 SSD Guard

SSD Guard, a feature that is unique to MegaRAID, increases the reliability of SSDs by automatically copying data from a drive with potential to fail to a designated hot spare or newly inserted drive. Because SSDs are more reliable than hard disk drives (HDDs), non-redundant RAID 0 configurations are much more common than in the past. SSD Guard offers added data protection for RAID 0 configurations.

SSD Guard works by looking for a predictive failure while monitoring the SDD Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) error log. If errors indicate that a SSD failure is imminent, the MegaRAID software starts a rebuild to preserve the data on the SSD and sends appropriate warning event notifications.

1.5 Dimmer Switch Features

Powering drives and cooling drives represent a major cost for data centers. The MegaRAID Dimmer Switch feature set reduces the power consumption of the devices connected to a MegaRAID controller. This helps to share resources more efficiently and lowers the cost.

Dimmer Switch I - Spin down unconfigured disks. This feature is configurable and can be disabled.

Dimmer Switch II - Spin down Hot Spares. This feature is configurable and can be disabled.

Dimmer Switch III - This new feature spins down any Logical Disk after 30 minutes of inactivity, by default, if the array can be spun up within 60 seconds. This feature is configurable and can be disabled.

1.6 UEFI 2.0 Support

UEFI 2.0 provides MegaRAID customers with expanded platform support. The MegaRAID UEFI 2.0 driver, a boot service device driver, handles block IO requests and SCSI pass-through (SPT) commands, and offers the ability to launch pre-boot MegaRAID management applications through a driver configuration protocol (DCP). The UEFI driver also supports driver diagnostic protocol, which allows administrators to access pre-boot diagnostics.

1.7 Configuration Scenarios

You can use the SAS RAID controllers in three scenarios:

- **Low-end, Internal SATA II Configurations**

In these configurations, use the RAID controller as a high-end SATA II-compatible controller that connects up to 8 disks either directly or through a port expander. These configurations are mostly for low-end or entry servers. Enclosure management is provided through out-of-band Inter-IC (I2C) bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.

- **Midrange Internal SAS Configurations**

These configurations are like the internal SATA II configurations, but with high-end disks. These configurations are more suitable for low-range to midrange servers.

- **High-end External SAS/SATA II Configurations**

These configurations are 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 shows a direct-connect configuration. The Inter-IC (I²C) 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.

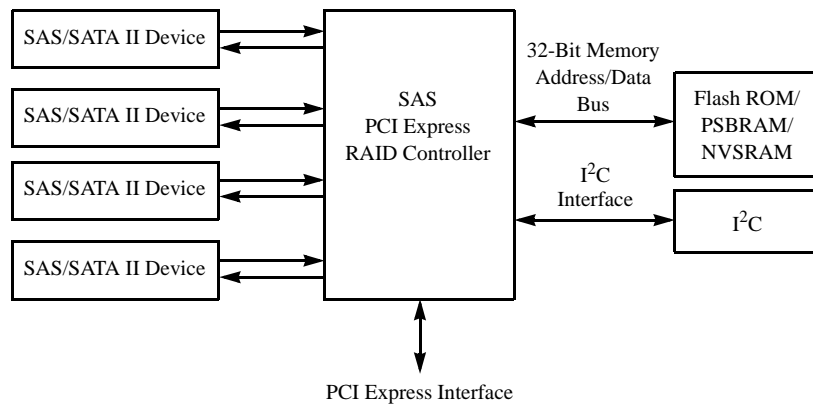


Figure 1: Example of an LSI SAS Direct-Connect Application

Figure 2 shows an example of a SAS RAID controller configured with an LSISASx12 expander that is connected to SAS disks, SATA II disks, or both.

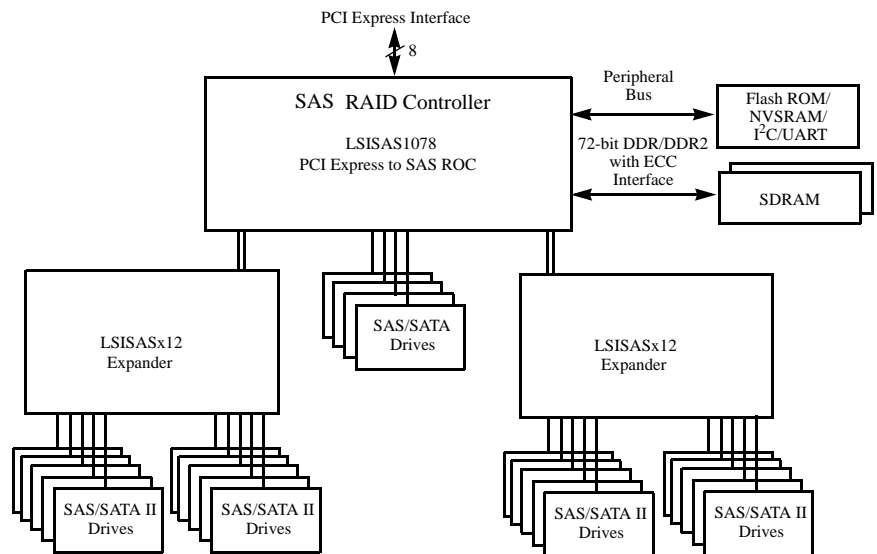


Figure 2: Example of an LSI SAS RAID Controller Configured with an LSISASx12 Expander

1.7.1 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 CacheCade 2.0 SSD Read Caching software devices.

You can choose whether to allow a virtual drive to consist of both CacheCade 2.0 SSD Read Caching software devices and HDDs. For a virtual drive that consists of CacheCade 2.0 SSD Read Caching software only, you can choose whether to allow SAS CacheCade 2.0 SSD Read Caching software drives and SATA CacheCade 2.0 SSD Read Caching software drives in that virtual drive. For virtual drives that have both CacheCade 2.0 SSD Read Caching software and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA CacheCade 2.0 SSD Read Caching software devices in various combinations.

[Table 1](#) lists the valid drive mix configurations you can use when you create virtual drives and allow HDD and CacheCade 2.0 SSD Read Caching software mixing. The valid drive mix configurations are based on manufacturer settings.

Table 1: Valid Drive Mix Configurations

#	Valid Drive Mix Configurations
1.	SAS HDD with SAS SDD (SAS-only configuration)
2.	SATA HDD with SATA CacheCade 2.0 SSD Read Caching software (SATA-only configuration)
3.	SAS HDD with a mix of SAS and SATA CacheCade 2.0 SSD Read Caching software (a SATA HDD cannot be added)
4.	SATA HDD with a mix of SAS and SATA CacheCade 2.0 SSD Read Caching software (a SAS HDD cannot be added)
5.	SAS CacheCade 2.0 SSD Read Caching software with a mix of SAS and SATA HDD (a SATA CacheCade 2.0 SSD Read Caching software cannot be added)
6.	SATA CacheCade 2.0 SSD Read Caching software with a mix of SAS and SATA HDD (a SAS CacheCade 2.0 SSD Read Caching software cannot be added)
7.	A mix of SAS and SATA HDD with a mix of SAS and SATA CacheCade 2.0 SSD Read Caching software
8.	A CacheCade 2.0 SSD Read Caching software cannot be added to a HDD, but a SAS/SATA mix is allowed.

NOTE: Only one of the valid configurations listed in [Table 1](#) is allowed based on your controller card manufacturing settings.

NOTE: The valid drive mix also applies to hot spares. For hot spare information, see [Section 2.4.12, Hot Spares, on page 28](#).

1.8 Technical Support

For assistance with installing, configuring, or running your MegaRAID 6Gb/s SAS RAID controllers, contact an LSI Technical Support representative.

Click the following link to access the LSI Technical Support page for storage and board support:

http://www.lsi.com/support/storage/tech_support/index.html

From this page, you can send an e-mail or call a Technical Support representative, or submit a new service request and view its status.

E-mail:

http://www.lsi.com/support/support_form.html

Phone Support:

http://www.lsi.com/support/storage/phone_tech_support/index.html

1-800-633-4545 (North America)

00-800-5745-6442 (International)

Chapter 2

Introduction to RAID

This chapter describes Redundant Array of Independent Disks (RAID), 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. Some common RAID functions follow:

- 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 on which to work

2.4 Components and Features

2.4.1 Drive Group

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.

A drive group is a group of physical drives. These 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 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 entire virtual drive (all drives associated with the virtual drive) fails.

Fault tolerance is often associated with system availability because it allows the system to be available during the failures. However, fault tolerance 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.

2.4.3.1 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 issues 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 ESMs, 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 shows 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.

2.4.4 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.

2.4.5 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 situation 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:

- 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.
- The rebuild operation takes precedence over the copyback operation when the conditions exist to start both operations. For example:
 - The hot spare is not configured (or unavailable) in the system.
 - Two drives (both members of virtual drives) exist, with one drive exceeding the SMART error threshold, and the other failed.
 - 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.

2.4.6 Background Initialization

Background initialization is a check for media errors on the drives when you create a virtual drive. It is an automatic operation that starts five minutes after you create the virtual drive. This check ensures that striped data segments are the same on all of the drives in the drive group.

Background initialization is similar to a consistency check. The difference between the two is that a background initialization is forced on new virtual drives and a consistency check is not.

New RAID 5 virtual drives and new RAID 6 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives, the background initialization does not start. The background initialization needs to be started manually. The following number of drives are required:

- New RAID 5 virtual drives must have at least five drives for background initialization to start.
- New RAID 6 virtual drives must have at least seven drives for background initialization to start.

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.

2.4.7 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 software to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See [Section 5.8, Controller Property-Related Options](#) and [Section 9.17, Running a Patrol Read](#).

2.4.8 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 3: Example of Disk Striping (RAID 0)

2.4.8.1 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.

2.4.8.2 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.

2.4.8.3 Strip Size

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

2.4.9 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 it is expensive because each drive in the system must be duplicated. [Figure 4](#) shows an example of disk mirroring.




Figure 4: Example of Disk Mirroring (RAID 1)


2.4.10 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](#).

Table 2: 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 5. 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.



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

Figure 5: Example of Distributed Parity (RAID 5)

2.4.11 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 6, 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 6: Example of Disk Spanning

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.

2.4.11.1 Spanning for RAID 00, RAID 10, RAID 50, and RAID 60

Table 3 describes how to configure RAID 00, RAID 10, RAID 50, and RAID 60 by spanning. The virtual drives must have the same stripe size and the maximum number of spans is 8. The full drive capacity is used when you span virtual drives; you cannot specify a smaller drive capacity.

See Chapter 8, *Configuration* for detailed procedures for configuring drive groups and virtual drives, and spanning the drives.

Table 3: 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 8 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.

NOTE: In a spanned virtual drive (R10, R50, R60) the span numbering starts from Span 0, Span 1, Span 2, and so on.

2.4.12 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.

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 after 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 drive failures are present on a split backplane configuration, the hot spare will be used first on the backplane side in which it resides.

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

NOTE: If a rebuild to a hot spare fails for any reason, the hot spare drive is marked as failed. If the source drive fails, both the source drive and the hot spare drive are marked as failed.

The hot spare can be of two types:

- Global hot spare
- Dedicated hot spare

2.4.12.1 Global Hot Spare

Use a global hot spare drive 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.

2.4.12.2 Dedicated Hot Spare

Use a dedicated hot spare 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 failover. 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 only to the same controller.
- 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 a 500-GB drive, the hot spare must be 500-GB or larger.

2.4.13 Disk Rebuilds

When a drive in a RAID drive group fails, you can rebuild the drive by re-creating the data that was stored on the drive before it failed. The RAID controller re-creates 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 resumes 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. (RAID-level migration changes a virtual drive from one RAID level to another.)

2.4.14 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 percent or 100 percent is not recommended. The default rebuild rate is 30 percent.

2.4.15 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 occurs automatically if these situation occurs:

- The newly inserted drive is the same capacity as or larger than the failed drive.
- The newly inserted drive 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.

2.4.16 Drive States

A drive state is a property indicating the status of the drive. The drive states are described in [Table 4](#).

Table 4: 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.
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.

Table 4: Drive States (Continued)

State	Description
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.

2.4.17 Virtual Drive States

The virtual drive states are described in [Table 5](#).

Table 5: 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.

2.4.18 Beep Codes

An alarm sounds on the MegaRAID controller when a virtual drive changes from an optimal state to another state, when a hot spare rebuilds, and for test purposes.

Table 6: Beep Codes, Events, and Virtual Drive States

Event	Virtual Drive State	Beep Code
RAID 0 virtual drive loses a virtual drives	Offline	3 seconds on and 1 second off
RAID 1 loses a mirror drive	Degraded	1 second on and 1 second off
RAID 1 loses both drives	Offline	3 seconds on and 1 second off
RAID 5 loses one drive	Degraded	1 second on and 1 second off
RAID 5 loses two or more drives	Offline	3 seconds on and 1 second off
RAID 6 loses one drive	Partially Degraded	1 second on and 1 second off
RAID 6 loses two drives	Degraded	1 second on and 1 second off
RAID 6 loses more than two drives	Offline	3 seconds on and 1 second off
A hot spare completes the rebuild process and is brought into a drive group	N/A	1 second on and 3 seconds off

2.4.19 Enclosure Management

Enclosure management is the intelligent monitoring of the disk subsystem by software, hardware or both. 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, the RAID controller 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.

RAID 1 uses mirroring so that data written to one drive is simultaneously written to another drive. RAID 1 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 any 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.

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

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 8 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 RAID 5 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. RAID 60 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 RAID 0 offers 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 entire virtual drive (all drives associated with the virtual drive) fails.

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 situation makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. [Table 7](#) provides an overview of RAID 0. [Figure 7](#) provides a graphic example of a RAID 0 drive group.

Table 7: RAID 0 Overview

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




Figure 7: RAID 0 Drive Group Example with Two Drives

2.5.4 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 through 32 in a single span. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. Table 8 provides an overview of RAID 1. Figure 8 provides a graphic example of a RAID 1 drive group.

Table 8: RAID 1 Overview

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




Figure 8: RAID 1 Drive Group

2.5.5 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 9 provides an overview of RAID 5. Figure 9 provides a graphic example of a RAID 5 drive group.

Table 9: RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to re-create all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak points	Not well-suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). Drive performance is 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.
Number of Drives in this RAID Level	3 through 32



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

Figure 9: RAID 5 Drive Group with Six Drives

2.5.6 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 any two drives in a virtual drive without losing data. RAID 6 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 re-create 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 10 provides an overview of a RAID 6 drive group.

Table 10: RAID 6 Overview

Uses	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.
Strong points	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.
Weak points	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.
Drives	3 through 32

Figure 10 shows a RAID 6 data layout. The second set of parity drives is denoted by Q. The P drives follow the RAID 5 parity scheme.



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

Figure 10: Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)

2.5.7 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 entire virtual drive (all drives associated with the virtual drive) fails.

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 situation makes RAID 00 ideal for applications that require high bandwidth but do not require fault tolerance. Table 11 provides an overview of RAID 00. Figure 11 provides a graphic example of a RAID 00 drive group.

Table 11: RAID 00 Overview

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




Figure 11: RAID 00 Drive Group Example with Two Drives

2.5.8 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1, and it 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 drive failures occur, 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 8 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 12 provides an overview of RAID 10.

Table 12: RAID 10 Overview

Uses	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.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	4 to 32 in multiples of 4 — The maximum number of drives supported by the controller (using an even number of drives in each RAID 10 virtual drive in the span).

In Figure 12, virtual drive 0 is created by distributing data across four drive groups (drive groups 0 through 3).




Figure 12: 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 8 spans and tolerate up to 8 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 13 provides an overview of RAID 50.

Table 13: RAID 50 Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium-to-large capacity.
Strong points	Provides high data throughput, data redundancy, and very good performance.
Weak points	Requires 2 times to 8 times as many parity drives as RAID 5.
Drives	8 spans of RAID 5 drive groups containing 3 to 32 drives each (limited by the maximum number of devices supported by the controller)




Figure 13: 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 14: RAID 60 Overview

Uses	<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 re-create 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>
Strong points	<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>
Weak points	<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>
Drives	A minimum of 8.

Figure 14 shows a RAID 60 data layout. The second set of parity drives is denoted by Q. The P drives follow the RAID 5 parity scheme.



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

Figure 14: RAID 60 Level Virtual Drive

2.6 RAID Configuration Strategies

The following factors in RAID drive group configuration are most important:

- 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 15: RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
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 performance 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. Because 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.
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.

Table 15: RAID Levels and Fault Tolerance (Continued)

RAID Level	Fault Tolerance
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 data 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 re-create all missing information. RAID 50 can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
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 re-create 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 16](#) describes the performance for each RAID level.

Table 16: 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.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.
5	RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Because each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware-based exclusive-or assist make RAID 5 performance exceptional in many different environments. 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.
6	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.
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.

Table 16: RAID Levels and Performance (Continued)

RAID Level	Performance
10	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 8.) 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.
50	RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the 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 8.) 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.
60	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 8.) 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. 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.

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 17](#) explains the effects of the RAID levels on storage capacity.

Table 17: 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. The usable capacity of a RAID 0 array is equal to the number of drives in the array into the capacity of the smallest drive in the array.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This situation is expensive because each drive in the system must be duplicated. The usable capacity of a RAID 1 array is equal to the capacity of the smaller of the two drives in the array.
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 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. The usable capacity of a RAID 5 array is equal to the number of drives in the array, minus one, into the capacity of the smallest drive in the array.
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 60 more expensive to implement. The usable capacity of a RAID 6 array is equal to the number of drives in the array, minus two, into the capacity of the smallest drive in the array.
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.

Table 17: RAID Levels and Capacity (Continued)

RAID Level	Capacity
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 large 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 situation 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 hot spares with enough capacity to rebuild the failed drives are not 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.

2.9 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.

2.9.1 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 18](#) 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 18: Factors to Consider for Drive Group Configuration

Requirement	Rank	Suggested RAID Levels
Storage space		RAID 0, RAID 5, RAID 00
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

SafeStore Disk Encryption

This chapter describes the LSI SafeStore™ Disk Encryption service. The SafeStore Disk Encryption service is a collection of features within LSI storage products that supports self-encrypting disks. SafeStore encryption services supports local key management.

3.1 Overview

The SafeStore Disk Encryption service offers the ability to encrypt data on drives and use disk-based key management to provide data security. This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting drives, if you remove a drive from its storage system or the server in which it is housed, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

With the SafeStore encryption service, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual disk (VD) level.

Any encryption solution requires management of the encryption keys. The security service provides a way to manage these keys. Both the WebBIOS Configuration Utility ([Chapter 4](#)) and the MegaRAID Storage Manager software ([Chapter 11](#)) offer procedures that you can use to manage the security settings for the drives.

3.2 Purpose and Benefits

Security is a growing market concern and requirement. MegaRAID customers are looking for a comprehensive storage encryption solution to protect data. You can use the SafeStore encryption service to help protect your data.

In addition, SafeStore local key management removes the administrator from most of the daily tasks of securing data, thereby reducing user error and decreasing the risk of data loss. Also, SafeStore local key management supports instant secure erase of drives that permanently removes data when repurposing or decommissioning drives. These services provide a much more secure level of data erasure than other common erasure methods, such as overwriting or degaussing.

3.3 Terminology

Table 19 describes the terminology related to the SafeStore encryption feature.

Table 19: Terminology Used in FDE

Option	Description
Authenticated Mode	The RAID configuration is keyed to a user password. The password must be provided on system boot to authenticate the user and facilitate unlocking the configuration for user access to the encrypted data.
Blob	A blob is created by encrypting a keys using another key. There are two types of blob in the system – encryption key blob and security key blob.
Key backup	You need to provide the controller with a lock key if the controller is replaced or if you choose to migrate secure virtual disks. To do this task, you must back up the security key.
Password	An optional authenticated mode is supported in which you must provide a password on each boot to make sure the system boots only if the user is authenticated. Firmware uses the user password to encrypt the security key in the security key blob stored on the controller.
Re-provisioning	Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For SafeStore encrypted drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted. This situation does not apply to controller-encrypted drives, because deleting the virtual disk destroys the encryption keys and causes a secure erase. See Section 3.5, Instant Secure Erase , for information about the instant secure erase feature.
Security Key	A key based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. If the security key is unavailable, user data is irretrievably lost. You must take all precautions to never lose the security key.
Un-Authenticated Mode	This mode allows controller to boot and unlock access to user configuration without user intervention. In this mode, the security key is encrypted into a security key blob, stored on the controller, but instead of a user password, an internal key specific to the controller is used to create the security key blob.
Volume Encryption Keys (VEK)	The controller uses the volume encryption keys to encrypt data when a controller-encrypted virtual disk is created. These keys are not available to the user. The firmware uses a unique 512-bit key for each virtual disk. The VEKs for the virtual disks are stored on the physical disks in a VEK blob.

3.4 Workflow

3.4.1 Enable Security

You can enable security on the controller. After you enable security, you have the option to create secure virtual drives using a security key.

There are three procedures you can perform to create secure virtual drives using a security key:

- Create the security key identifier
- Create the security key
- Create a password (optional)

3.4.1.1 Create the Security Key Identifier

The security key identifier appears whenever you enter the security key. If you have multiple security keys, the identifier helps you determine which security key to enter. The controller provides a default identifier for you. You can use the default setting or enter your own identifier.

3.4.1.2 Create the Security Key

You need to enter the security key to perform certain operations. You can choose a strong security key that the controller suggests.

CAUTION: If you forget the security key, you will lose access to your data.

3.4.1.3 Create a Password

The password provides additional security. The password should be different from the security key. You can select a setting in the utilities so that you must enter the password whenever you boot your server.

CAUTION: If you forget the password, you will lose access to your data.

When you use the specified security key identifier, security key, and password, security is enabled on the controller.

3.4.2 Change Security

You can change the security settings on the controller, and you have the option to change the security key identifier, security key, and password. If you have previously removed any secured drives, you still need to supply the old security key to import them.

You can perform three procedures to change the security settings on the controller:

- Change the security key identifier
- Change the security key
- Change a password

See [Section 4.7, *Selecting SafeStore Encryption Services Security Options*](#) for the procedures used to change security options in WebBIOS or [Section 11.6, *LSI SafeStore Encryption Services*](#) for the procedures used to change security options in the MegaRAID Storage Manager software.

3.4.2.1 Change the Security Key Identifier

You have the option to edit the security key identifier. If you plan to change the security key, it is highly recommended that you change the security key identifier. Otherwise, you will not be able to differentiate between the security keys.

You can select whether you want to keep the current security key identifier or enter a new one. To change the security key identifier, enter a new security key identifier.

3.4.2.2 Change the Security Key

You can choose to keep the current security key or enter a new one. To change the security key, you can either enter the new security key or accept the security key that the controller suggests.

3.4.2.3 Add or Change the Password

You have the option to add a password or change the existing one. To change the password, enter the new password. To keep the existing password, enter the current password. If you choose this option, you must enter the password whenever you boot your server.

This procedure updates the existing configuration on the controller to use the new security settings.

3.4.3 Create Secure Virtual Drives

You can create a secure virtual drive and set its parameters as desired. To create a secure virtual drive, select a configuration method. You can select either simple configuration or advanced configuration.

3.4.3.1 Simple Configuration

If you select simple configuration, select the redundancy type and drive security method to use for the drive group.

See [Section 8.1.4, *Creating a Virtual Drive Using Simple Configuration*](#) for the procedures used to select the redundancy type and drive security method for a configuration.

3.4.3.2 Advanced Configuration

If you select advanced configuration, select the drive security method, and add the drives to the drive group.

See [Section 8.1.5, *Creating a Virtual Drive using Advanced Configuration*](#) for the procedures used to import a foreign configuration.

After the drive group is secured, you cannot remove the security without deleting the virtual drives.

3.4.4 Import a Foreign Configuration

After you create a security key, you can run a scan for a foreign configuration and import a locked configuration. (You can import unsecured or unlocked configurations when security is disabled.) A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. WebBIOS Configuration Utility and the MegaRAID Storage Manager software allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See [Section 4.8, *Viewing and Changing Device Properties*](#) for the procedure used to import a foreign configuration in WebBIOS or [Section 11.6.12, *Importing or Clearing a Foreign Configuration*](#) for the procedure in the MegaRAID Storage Manager software.

To import a foreign configuration, you must first enable security to allow importation of locked foreign drives. If the drives are locked and the controller security is disabled, you cannot import the foreign drives. Only unlocked drives can be imported when security is disabled.

After you enable the security, you can import the locked drives. To import the locked drives, you must provide the security key used to secure them. Verify whether any drives are left to import as the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all of the drives are imported, there is no configuration to import.

3.5 Instant Secure Erase

Instant Secure Erase is a feature used to erase data from encrypted drives. After the initial investment for an encrypted disk, there is no additional cost in dollars or time to erase data using the Instant Secure Erase feature.

You can change the encryption key for all MegaRAID RAID controllers that are connected to encrypted drives. All encrypted drives, whether locked or unlocked, always have an encryption key. This key is set by the drive and is always active. When the drive is unlocked, the data to host from the drive (on reads) and from the host to the drive cache (on writes) is always provided. However, when resting on the drive platters, the data is always encrypted by the drive.

You might not want to lock your drives because you have to manage a password if they are locked. Even if you do not lock the drives, there is still a benefit to using encrypted disks.

If you are concerned about data theft or other security issues, you might already invest in drive disposal costs, and there are benefits to using SafeStore encryption over other technologies that exist today, both in terms of the security provided and time saved.

If the encryption key on the drive changes, the drive cannot decrypt the data on the platters, effectively erasing the data on the disks. The National Institute of Standards and Technology (<http://www.nist.gov>) values this type of data erasure above secure erase and below physical destruction of the device.

Consider the following reasons for using instant secure erase.

If you need to repurpose the hard drive for a different application.

You might need to move the drive to another server to expand storage elsewhere, but the drive is in use. The data on the drive might contain sensitive data including customer information that, if lost or divulged, could cause an embarrassing disclosure of a security hole. You can use the instant secure erase feature to effectively erase the data so that the drive can be moved to another server or area without concern that old data could be found.

If you need to replace drives.

If the amount of data has outgrown the storage system, and there is no room to expand capacity by adding drives, you might choose to purchase upgrade drives. If the older drives support encryption, you can erase the data instantly so the new drives can be used.

If you need to return a disk for warranty activity.

If the drive is beginning to show SMART predictive failure alerts, you might want to return the drive for replacement. If so, the drive must be effectively erased if there is sensitive data. Occasionally a drive is in such bad condition that standard erasure applications do not work. If the drive still allows any access, it might be possible to destroy the encryption key.

Chapter 4

WebBIOS Configuration Utility

This chapter describes the WebBIOS Configuration Utility (CU), which enables you to create and manage RAID configurations on LSI SAS controllers.

4.1 Overview

The WebBIOS configuration utility, unlike the MegaRAID Storage Manager software, resides in the SAS controller BIOS and operates independently of the operating system.

You can use the WebBIOS configuration utility to perform the following tasks:

- Create drive groups and virtual drives for storage configurations.
- Display controller, drive, virtual 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.
- Create a CacheCade 2.0 SSD Read Caching configuration.

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

4.2 Starting the WebBIOS configuration utility

To start the WebBIOS configuration utility, perform the following steps:

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

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

The Controller Selection dialog appears.

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

The main WebBIOS configuration utility dialog appears.

NOTE: On systems that do not have the PS2 port, you must enable 'port 60/64 emulation' in the System BIOS to emulate USB as PS2. When this option is disabled on this system, WebBIOS does not work.

4.3 WebBIOS configuration utility Main Dialog Options

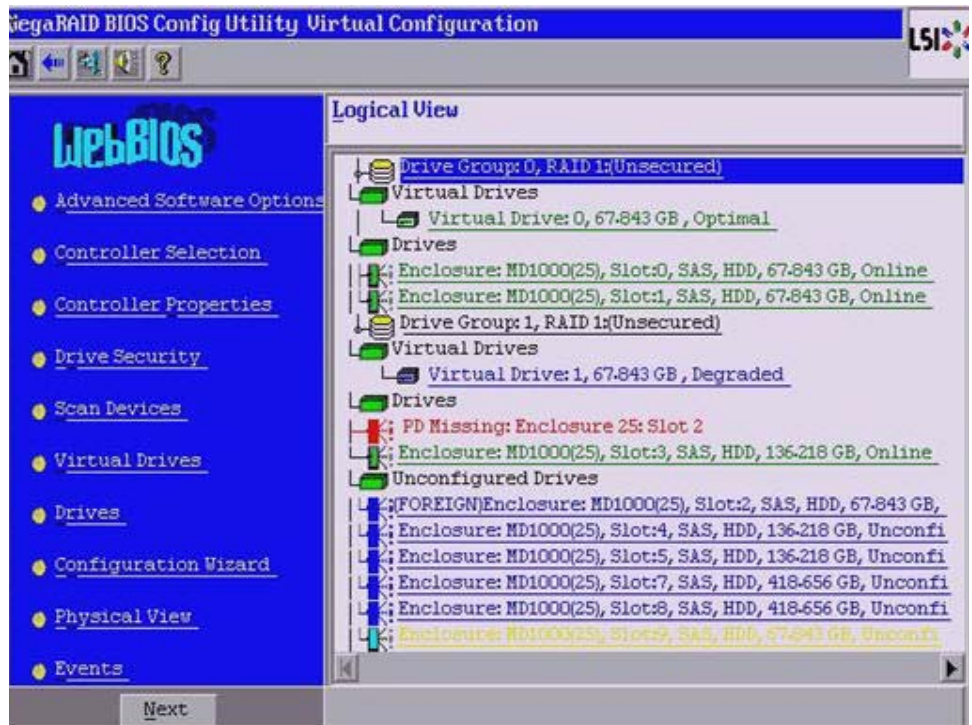


Figure 15: WebBIOS Configuration Utility Main Dialog

In the right frame, the dialog shows the virtual drives configured on the controller, and the drives that are connected to the controller. In addition, the dialog 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 dialog resolution for WebBIOS is 640 x 480.

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






NOTE: Unconfigured Bad drives are only displayed in the Physical View.

For drives in an enclosure, the dialog 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 configuration utility has the following buttons, as listed in [Table 20](#).

Table 20: WebBIOS configuration utility Toolbar Icons

Icon	Description
	Click this icon to return to the main dialog from any other WebBIOS configuration utility dialog.
	Click this icon to return to the previous dialog that you were viewing.
	Click this icon to exit the WebBIOS configuration utility wizard.
	Click this icon to turn off the sound on the onboard controller alarm.
	Click this icon to display information about the WebBIOS configuration utility version, bus number, and device number.

The following is a description of the options listed on the left frame of the WebBIOS configuration utility main dialog (the hotkey shortcut for each option is shown in parentheses next to the option name):

- **Advanced Software Options: (Alt+a)** Select this option to enable the advanced features in the controller. For more information, see [Section 4.4.1, Managing MegaRAID Advanced Software Options](#).
- **Controller Selection: (Alt+c)** Select this option to view the Controller Selection dialog, where you can select a different SAS controller. You can also view information about the controller and the devices connected to it, or create a new configuration on the controller.
- **Controller Properties: (Alt+p)** Select this option to view the properties of the currently selected SAS controller. For more information, see [Section 4.8.1, Viewing Controller Properties](#).

- **Drive Security: (Alt+r)** Select this option to encrypt data on the drives and use disk-based key management for the data security solution. This solution protects your data in case of theft or loss of physical drives. For more information, see [Section 4.7, *Selecting SafeStore Encryption Services Security Options*](#).
- **Scan Devices: (Alt+s)** Select this option to have the WebBIOS configuration utility re-scan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS configuration utility displays the results of the scan in the physical and virtual drive descriptions.
- **Virtual Drives: (Alt+v)** Select this option to view the Virtual Drives dialog, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see [Section 4.8.2, *Viewing Virtual Drive Properties, Policies, and Operations*](#).
- **Drives: (Alt+d)** Select this option to view the Drives dialog, where you can view drive properties, create hot spares, and perform other tasks. For more information, see [Section 4.8.3, *Viewing Drive Properties*](#).
- **Configuration Wizard: (Alt+o)** 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 4.5, *Creating a Storage Configuration*](#).
- **Logical View/Physical View: (Alt+l)** for the Logical view; **Alt+h** for the Physical view) Select this option to toggle between the Physical View dialog and the Logical View dialog.
- **Events: (Alt+e)** Select this option to view system events in the Event Information dialog. For more information, see [Section 4.13, *Viewing System Event Information*](#).
- **Exit: (Alt+x)** Select this option to exit the WebBIOS configuration utility and continue with system boot.

4.4 Managing Software Licensing

The MegaRAID advanced software offers the software license key feature to enable the advanced options in WebBIOS. The license key, also known as the *Activation key* is used to transfer the advanced features from one controller to another by configuring the Key Vault.

You need to configure the Advanced Software Options menu present in the WebBIOS main dialog to use the advanced features present in the controller.

4.4.1 Managing MegaRAID Advanced Software Options

Perform the following steps to configure the Advanced Software Options wizard to enable the advanced options using the activation key.

1. Click **Advanced Software Options** menu on the WebBIOS main dialog.

The Advanced Software Options wizard appears, as shown in the following figure.

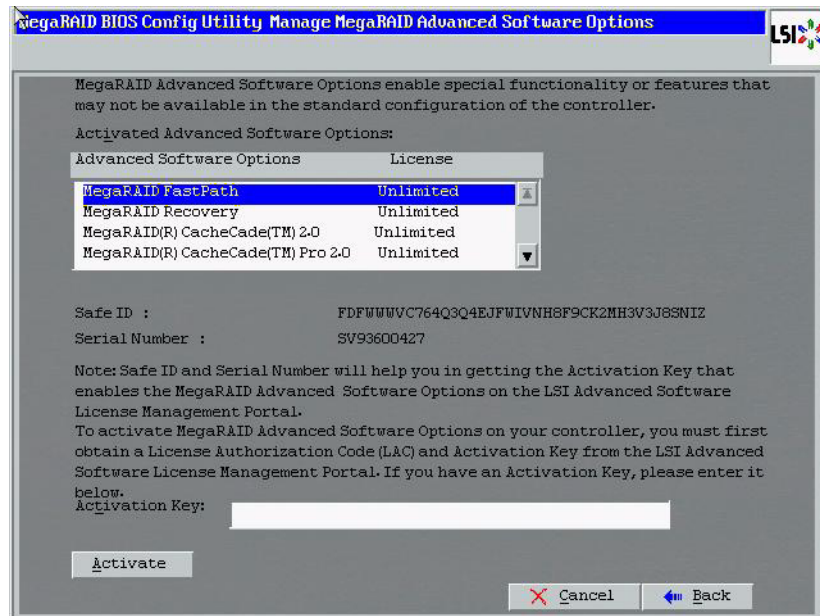


Figure 16: Manage MegaRAID Advanced Software Options Wizard

NOTE: When you click the **Advanced Software Options** menu in the main WebBIOS dialog, if re-hosting is not required, the Manage MegaRAID Advanced Software Options dialog appears; otherwise, if the user decides to opt for the re-hosting process, the [Section 4.4.8, Confirm Re-hosting Process](#) dialog appears.

The **Activated Advanced Software Options** field consists of **Advanced Software Options**, **License**, and **Mode** columns.

- The **Advanced Software** option column displays the list of advanced software features available in the controller.
- The **License** column displays the license details for the list of advanced softwares present in the **Advanced Software** options column. The license details validates if the software is under trial period, or if it can be used without any trial period (Unlimited).
- The **Mode** column displays the current status of the advanced software. The current status can be Secured, Not secured, or Factory installed.

Both the **Safe ID** and the **Serial Number** fields consist of a pre-defined value internally generated by the controller.

2. Click **Activate**.

The Advanced Software Options Summary wizard appears, as shown in [Figure 22](#).

3. Click **Configure Key Vault**.

The Confirm Rehosting Process wizard appears, as shown in [Figure 27](#).

The **Configure Key Vault** button is conditional, and appears in two scenarios.

- Scenario #1
When features have been transferred from NVRAM to key vault, and no re-hosting is required, the **Configure Key Vault** button is not displayed.
 - Scenario# 2
When the re-hosting process needs to be completed, the **Configure Key Vault** button appears.
4. Click **Deactivate All Trial Software**.
- The WebBIOS Deactivate All Trial Advanced Software Options dialog appears, as shown in the following figure.

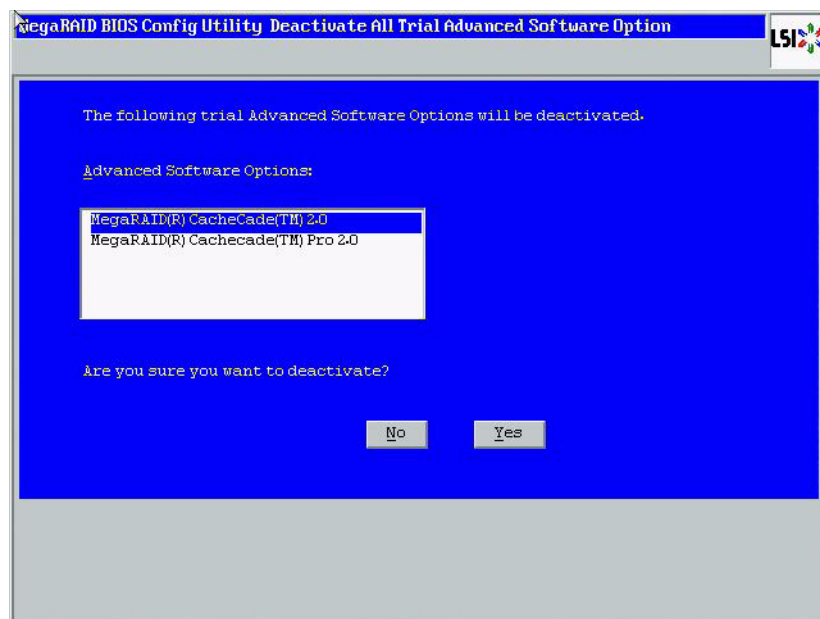


Figure 17: Deactivate All Trial Advanced Software Options Dialog

To *deactivate* the software that is being used with a trial key, click **Yes**; otherwise, click **No**.

When the activation key is improper in the **Activation** field in the Advanced Software Options wizard, the following messages appear based on the scenarios.

- Scenario # 1
If you enter an *invalid* activation key, the following message appears.



Figure 18: Invalid Activation Key Message

— Scenario # 2

If you leave the activation key field *blank* or enter *space* characters, the following message appears.




Figure 19: Config Utility Cannot Activate Advanced Software Options Message

— Scenario # 3

If you enter an *incorrect* activation key, and if there is a mismatch between the activation key and the controller, the following message appears.



Figure 20: Activation Key Mismatch Message

4.4.2 Reusing the Activation Key

If you are using an existing activated key, the features are transferred to the key vault, and the message appears as shown in the following figure.




Figure 21: Reusing the Activation Key

4.4.3 Managing Advanced Software Summary

When you click **Activate** in Manage MegaRAID Advanced Software Options dialog, the Advanced Software Options Summary wizard appears, as shown in the following figure.

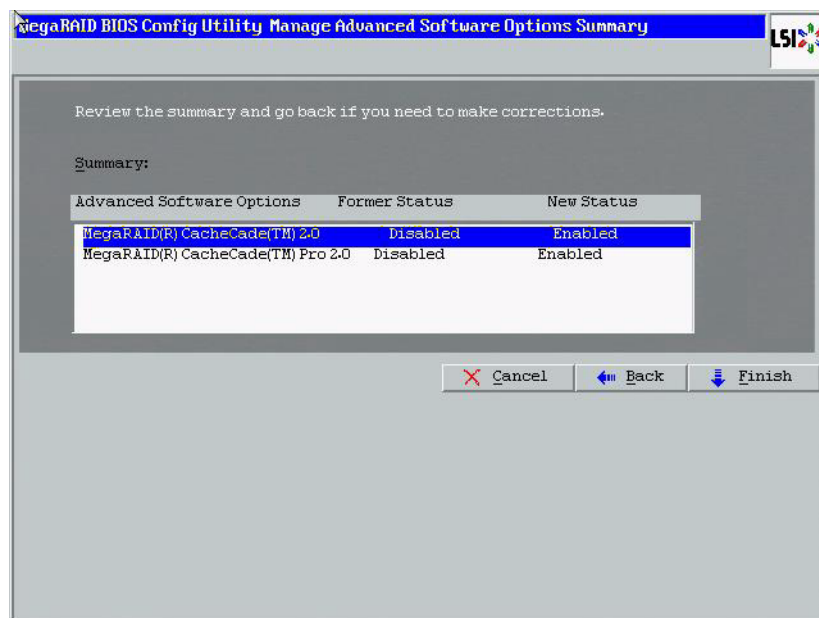


Figure 22: Advanced Software Options Summary Wizard

The **Summary** field displays the list of the advanced software options along with their *former status* and *new status* in the controller.

- The **Advanced Software Options** column displays the currently available software in the controller.
- The **Former Status** column displays the status of the available advanced software prior to entering the activation key.
- The **New Status** column displays the status of the available advanced software, after entering the activation key.

4.4.4 Activating an Unlimited Key Over a Trial Key

When you activate an unlimited key over a trial key, the Review the summary and go back if you need to make corrections message appears, as shown in the following figure.

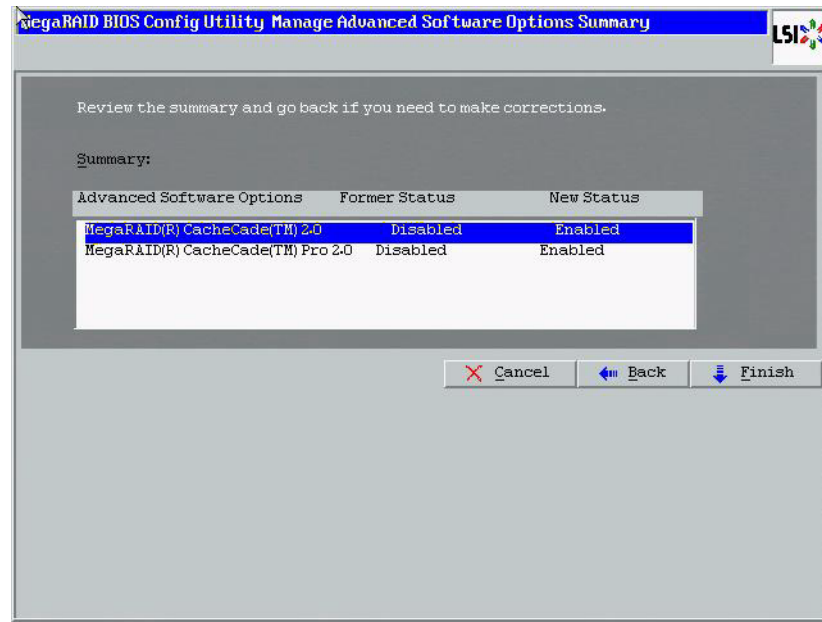


Figure 23: Activating an Unlimited Key over a Trial Key

4.4.5 Activating a Trial Software

When you activate a trial software, the This trial software expires in 30 days message appears, as shown in the following figure.

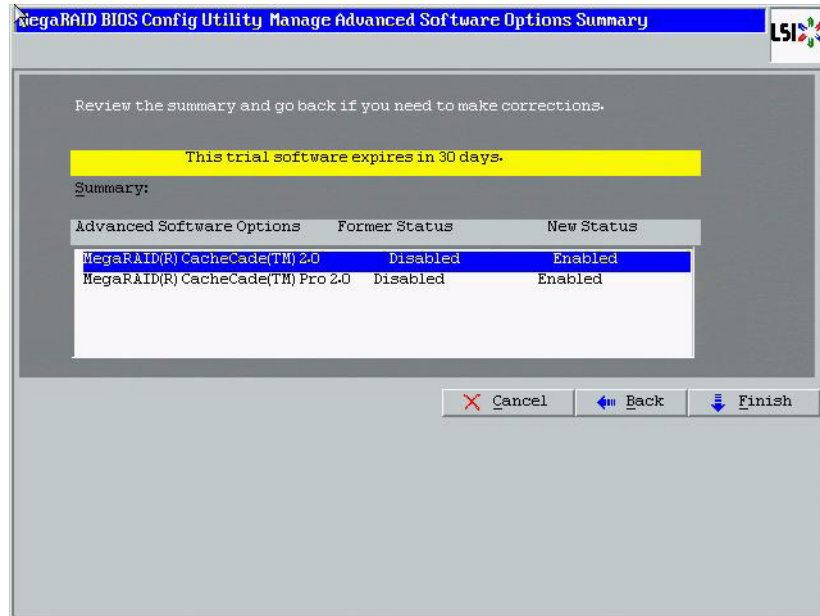


Figure 24: Activating a Trial Software Application

4.4.6 Activating an Unlimited Key

When you activate an unlimited key, the Review the summary and go back if you need to make corrections message appears, as shown in the following figure.

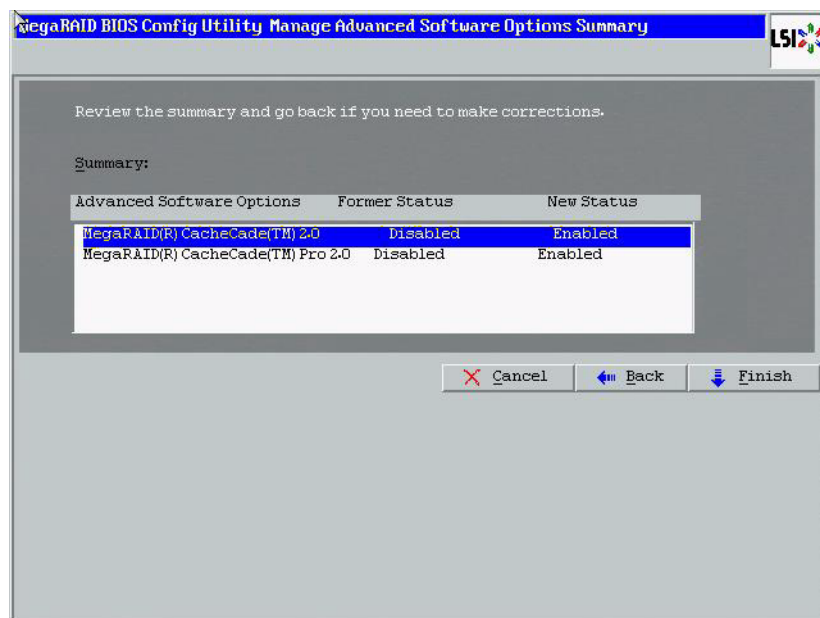


Figure 25: Activating an Unlimited Key

4.4.7 Securing MegaRAID Advanced Software

If the advanced software is not secured, when you click the **Configure Key Vault** button in the Advanced Software Options wizard, the WebBIOS Secure MegaRAID Advanced Software Options dialog box appears, as shown in the following figure.



Figure 26: Secure Advanced Software Options

4.4.8 Confirm Re-hosting Process

The confirming re-hosting process involves the process of transferring or re-hosting the advanced software features from one controller to another.

When you need to transfer the features from one controller (example, controller 1) to another controller (example, controller 2) and in the controller 2 NVRAM, if there are some features that need to be transferred to key vault, the Confirm Re-hosting Process dialog appears as shown in [Figure 27](#).

Perform the following steps to confirm the rehosting process.

1. Click the **Configure Key Vault** button in the **Advanced Software Options** wizard. The **Confirm Rehosting Process** wizard appears as shown in the following figure.

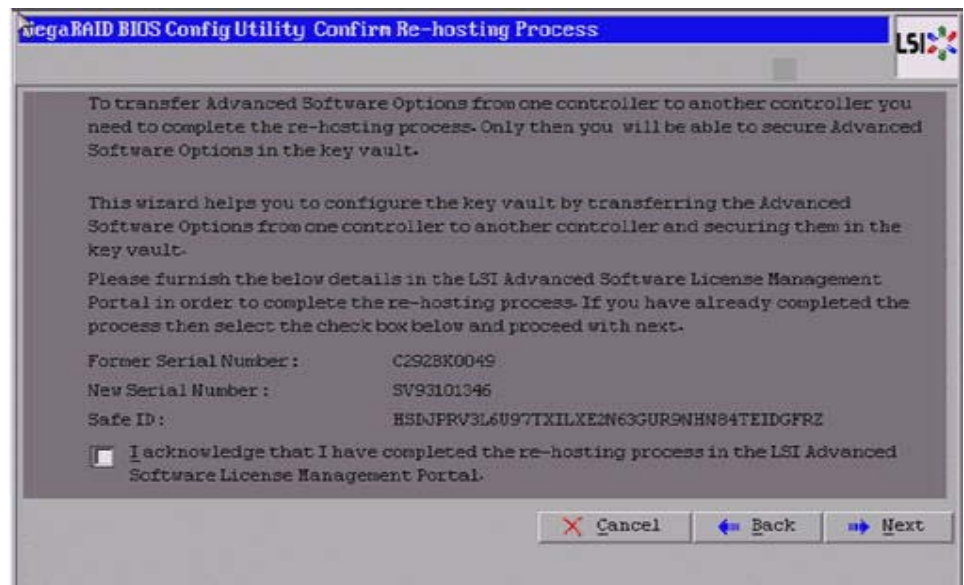


Figure 27: Confirm Re-hosting Process Dialog

2. Select the **I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal** check box.
3. Click **Next**.

The Manage Advanced Software Options Summary dialog appears as shown in [Figure 23](#).

4.4.9 Re-hosting Process Complete

In a scenario where only key vault feature needs to be transferred from controller 1 to controller 2, the Re-hosting Process Complete dialog appears as shown in the following figure.




Figure 28: Re-hosting Process Complete Dialog

1. Select the **I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal** check box.
2. Click **Next**.

The **Manage MegaRAID Advanced Software Options** Wizard appears.

The rehosting process is completed.

NOTE: If you click **Next** in the **Re-hosting Process Complete** dialog, if re-hosting is not complete, the features are not copied into the key vault, and the features remain in the key vault itself, but you can still use the advanced features.

4.5 Creating a Storage Configuration

This section explains how to use the WebBIOS configuration utility Configuration Wizard to configure RAID drive groups and virtual drives to create storage configurations:

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

1. Click **Configuration Wizard** on the WebBIOS main dialog.

The first Configuration Wizard dialog appears, as shown in the following figure.

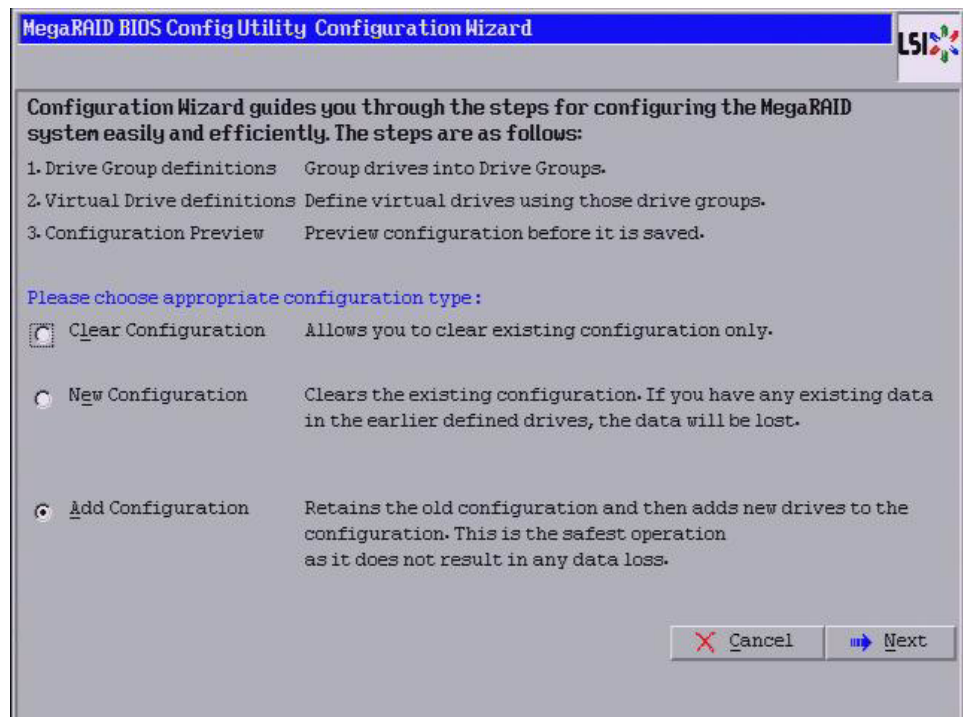


Figure 29: WebBIOS Configuration Wizard Dialog

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 option 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 Convert JBOD Drives to Unconfigured Drives dialog appears, as shown in the following figure.

NOTE: The JBOD Drives to Unconfigured Drives dialog appears only if the system detects JBOD drives.




Figure 30: JBOD Drives to Unconfigured Good Dialog

- Click **Next**.

The **WebBIOS Configuration Method** dialog appears, as shown in the following figure.

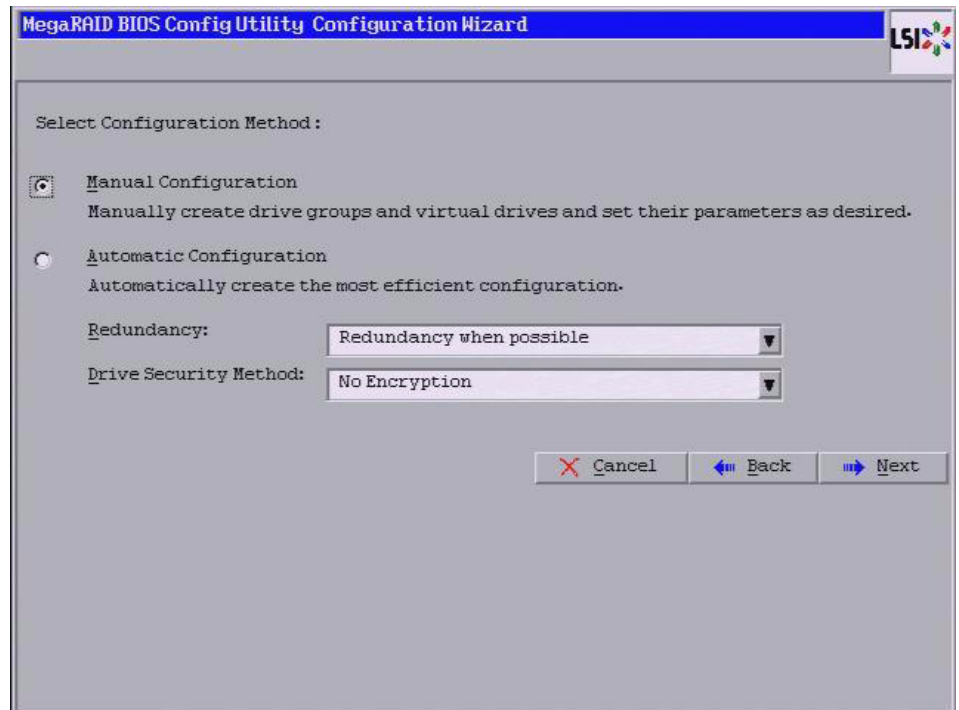


Figure 31: WebBIOS Configuration Method Wizard

6. 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.

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** drop down list:

- **Redundancy when possible**
- **No redundancy**

If you select **Automatic Configuration**, you can choose whether to use a drive security method. Select one of the following options in the **Drive Security Method** drop down list:

- **No Encryption**
- **Drive Encryption**

7. Click **Next** to continue.

If you select the **Automatic Configuration** radio button, continue with [Section 4.5.1, Using Automatic Configuration](#). If you select **Manual Configuration**, continue with [Section 4.5.2, Using Manual Configuration](#).

4.5.1 Using Automatic Configuration

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

1. When WebBIOS displays the proposed new configuration, review the information on the dialog, 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 drives.

WebBIOS configuration utility begins a background initialization of the virtual drives.

New RAID 5 virtual drives and new RAID 6 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives, the background initialization will not start. The following number of drives is required:

- New RAID 5 virtual drives must have at least five drives for a background initialization to start.
- New RAID 6 virtual drives must have at least seven drives for a background initialization to start.

4.5.2 Using Manual Configuration

This section contains the procedures for creating RAID drive groups for RAID levels 0, 1, 5, 6, 00, 10, 50, and 60.

4.5.2.1 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 dialog appears. Use this dialog 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 **Drive Groups** panel on the right, as shown in [Figure 32](#).

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

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

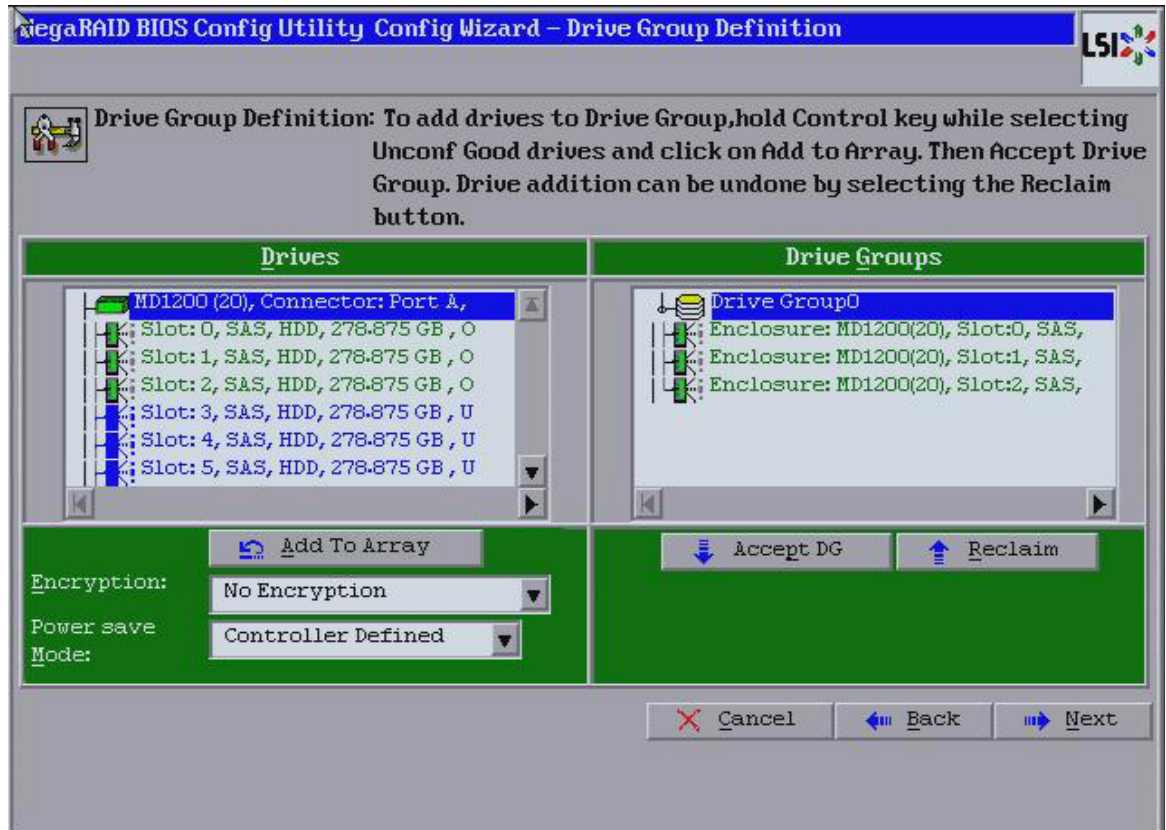


Figure 32: Drive Group Definition Dialog

5. After you finish selecting drives for the drive group, click **Accept DG**.
6. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 33](#). This dialog lists the possible RAID levels for the drive group.

Use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

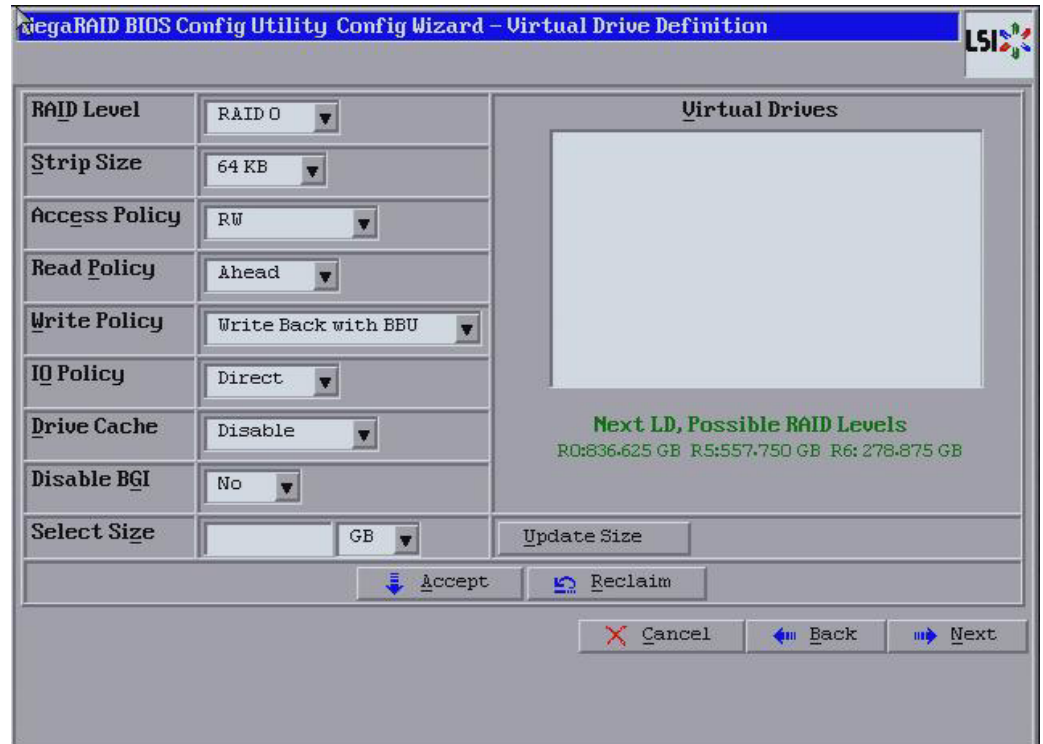


Figure 33: WebBIOS Virtual Drive Definition Dialog

7. Change the virtual drive options from the defaults listed on the dialog as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down list shows the possible RAID levels for the virtual drive. Select **RAID 0**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is **64 KB**.
- **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 option disables the read ahead capability. This option 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This option is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. 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 a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status.
 - No:** Leave background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to perform other configuration tasks. This option is the default setting.
 - 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 MB, GB, or TB. Usually, this is the full size for RAID 0 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size value for the selected RAID levels
8. Click **Accept** to accept the changes to the virtual drive definition.
If you need to undo the changes, click **Reclaim**.
 9. Click **Next** after you finish defining the virtual drives.
The Configuration Preview dialog appears, as shown in [Figure 34](#).

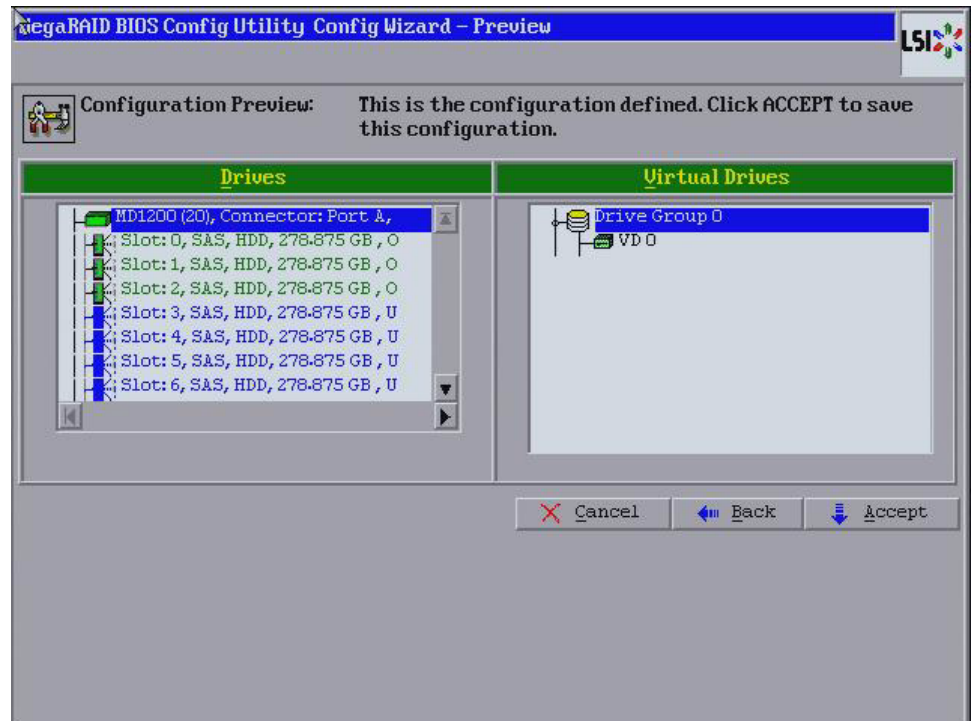


Figure 34: RAID 0 Configuration Preview Dialog

10. Check the information in the Configuration Preview Dialog.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.
The WebBIOS main menu appears.

4.5.2.2 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 dialog appears. Use this dialog 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 35](#).

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

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

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.

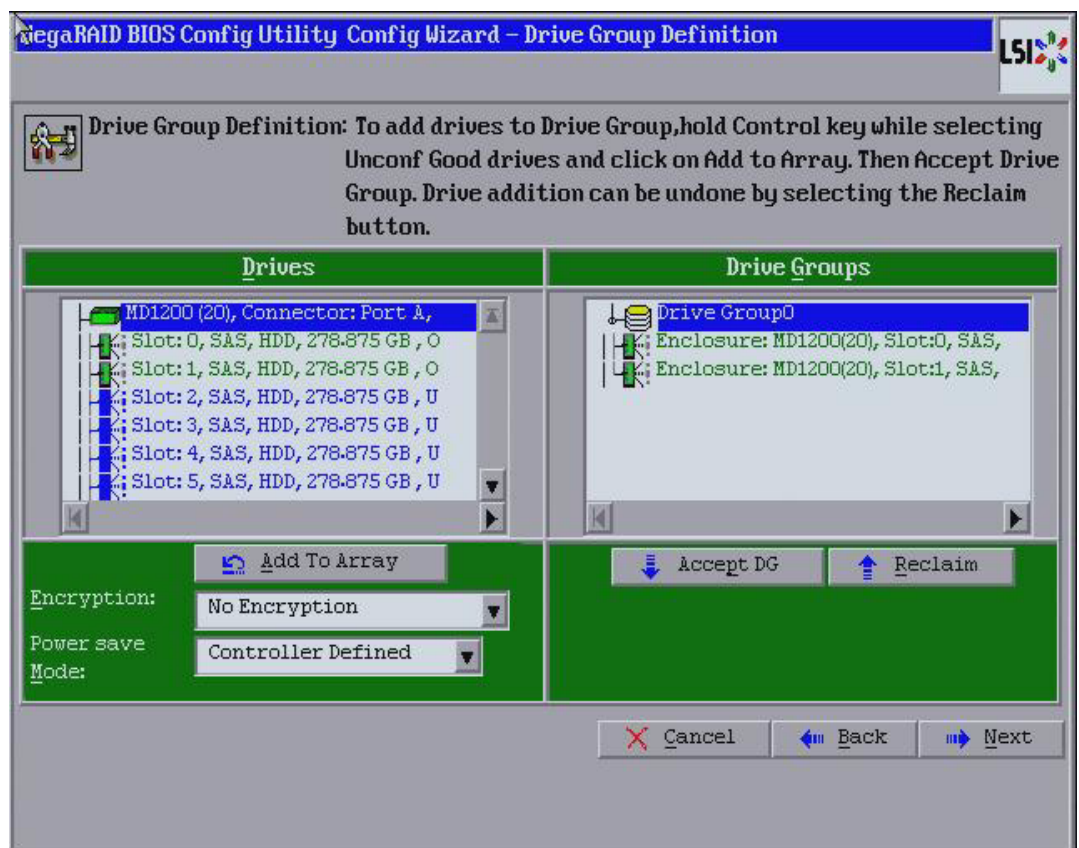


Figure 35: Drive Group Definition Dialog

5. After you finish selecting drives for the drive group, click **Accept DG**.
6. Click **Next**.

The Drive Group Definition dialog appears, as shown in [Figure 36](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

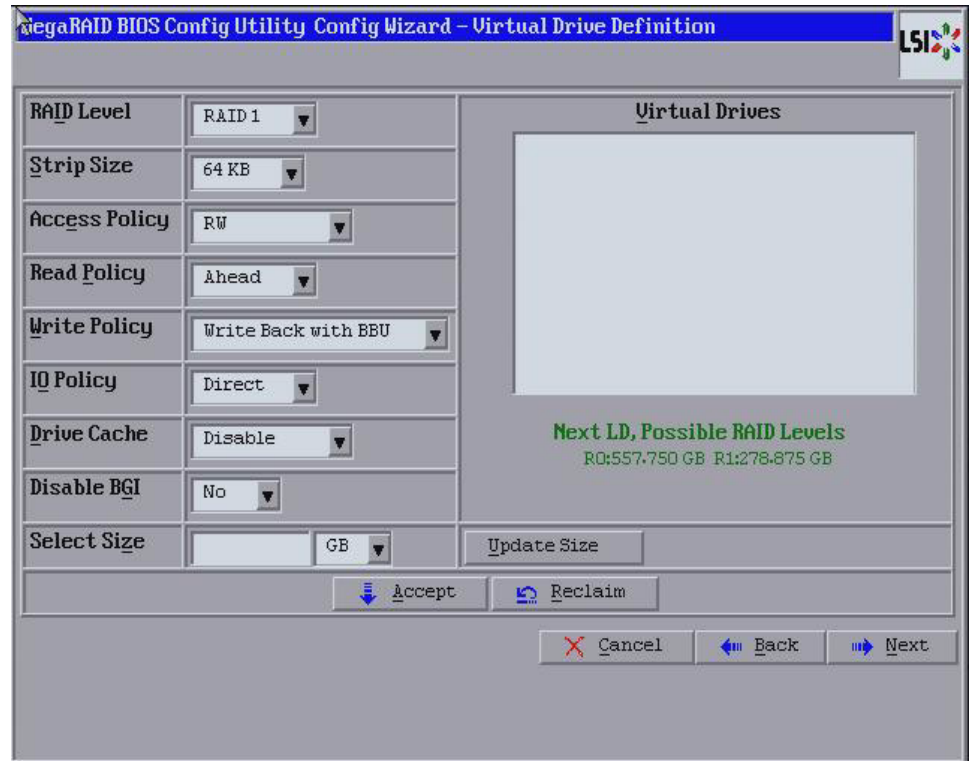


Figure 36: Virtual Group Drive Definition Dialog

7. Change the virtual drive options from the defaults listed on the dialog as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down list shows the possible RAID levels for the virtual drive. Select **RAID 1**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is **64 KB**.
- **Access Policy:** Select the type of data access that is allowed for this virtual drive.
 - RW:** Allow read/write access. This option is the default setting.
 - Read Only:** Allow read-only access.
 - Blocked:** Do not allow access.

- **Read Policy:** Specify the read policy for this virtual drive
 - Normal:** This option disables the read ahead capability. This option is the default setting.
 - 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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.
This is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. LSI recommends that you use *either* a battery to protect the controller cache, or an UPS to protect the entire system. If you do not use a battery or a UPS, and a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status:
 - No:** Leave background initialization enabled, which 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 drives in MB, GB, or TB. Usually, this would be the full size for RAID 1 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels
8. Click **Accept** to accept the changes to the virtual drive definition.
If you need to undo the changes, click **Reclaim**.
 9. Click **Next** after you finish defining the virtual drives.
The Configuration Preview dialog appears, as shown in [Figure 37](#).

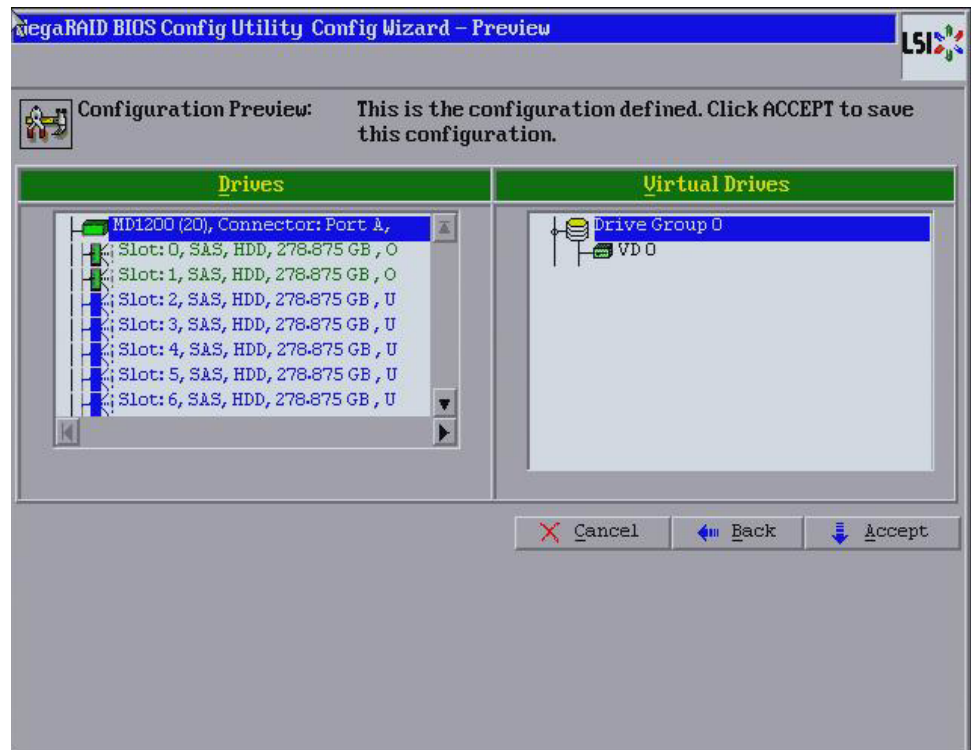


Figure 37: RAID 1 Configuration Preview Dialog

10. Check the information in the Configuration Preview dialog.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.
The WebBIOS main menu appears.

4.5.2.3 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 re-create 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 dialog appears. You use this dialog 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 38](#).
3. If you need to undo the changes, click **Reclaim**.
4. Choose whether to use power save mode.
5. Choose whether to use drive encryption.

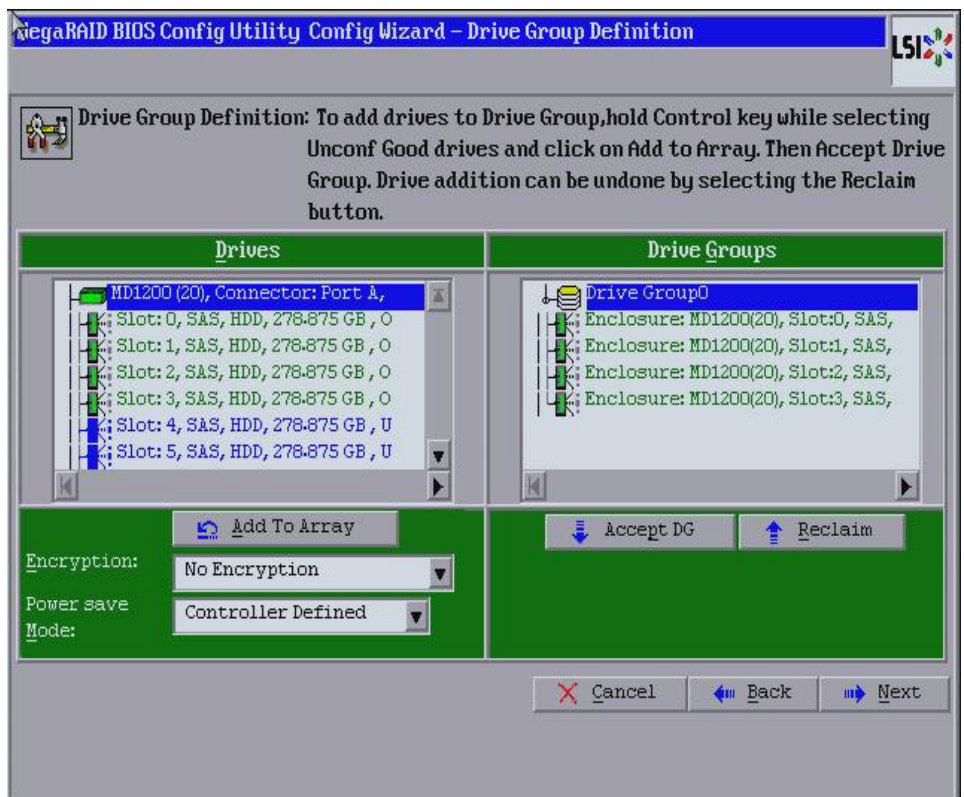


Figure 38: Drive Group Definition Dialog

6. After you finish selecting drives for the drive group, click **Accept DG**.
7. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 39](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

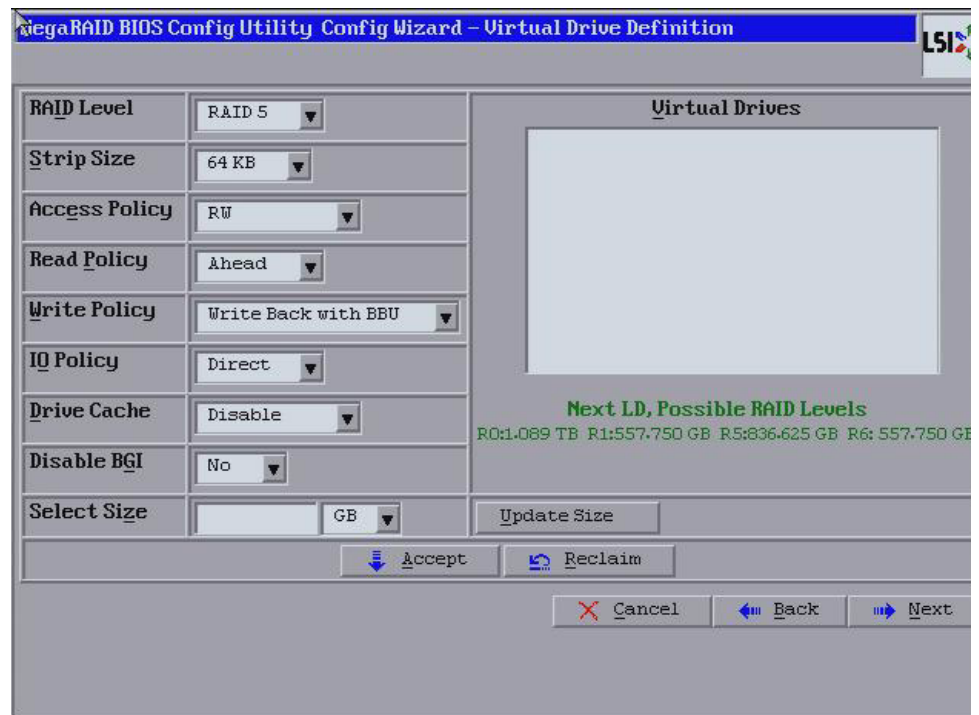


Figure 39: Virtual Drive Definition Dialog

8. Change the virtual drive options from the defaults listed on the dialog as needed.

Here are brief explanations of the virtual drive options.

- **RAID Level:** The drop-down list provides the possible RAID levels for the virtual drive. Select **RAID 5**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is 64 KB.
- **Access Policy:** Select the type of data access that is allowed for this virtual drive.
 - RW:** Allow read/write access. This option is the default setting.
 - Read Only:** Allow read-only access.
 - Blocked:** Do not allow access.

- **Read Policy:** Specify the read policy for this virtual drive.
 - Normal:** This option disables the read ahead capability. This option is the default setting.
 - 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This option is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. LSI recommends that you use *either* a battery to protect the controller cache, or a UPS to protect the entire system. If you do not use a battery or a UPS, and a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status.
 - No:** Leave background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to perform other configuration tasks. This option is the default setting.
 - Yes:** Select **Yes** if you do not want to allow background initializations for configurations on this controller.

NOTE: New RAID 5 virtual drives require at least five drives for a background initialization to start.

- **Select Size:** Specify the size of the virtual drive in MB, GB, or TB. Usually, this setting would be the full size for RAID 5 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.
- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels

Click **Accept** to accept the changes to the virtual drive definition.

9. If you need to undo the changes, click **Reclaim**.
10. Click **Next** after you finish defining the virtual drives.

The Configuration Preview dialog appears, as shown in [Figure 40](#).

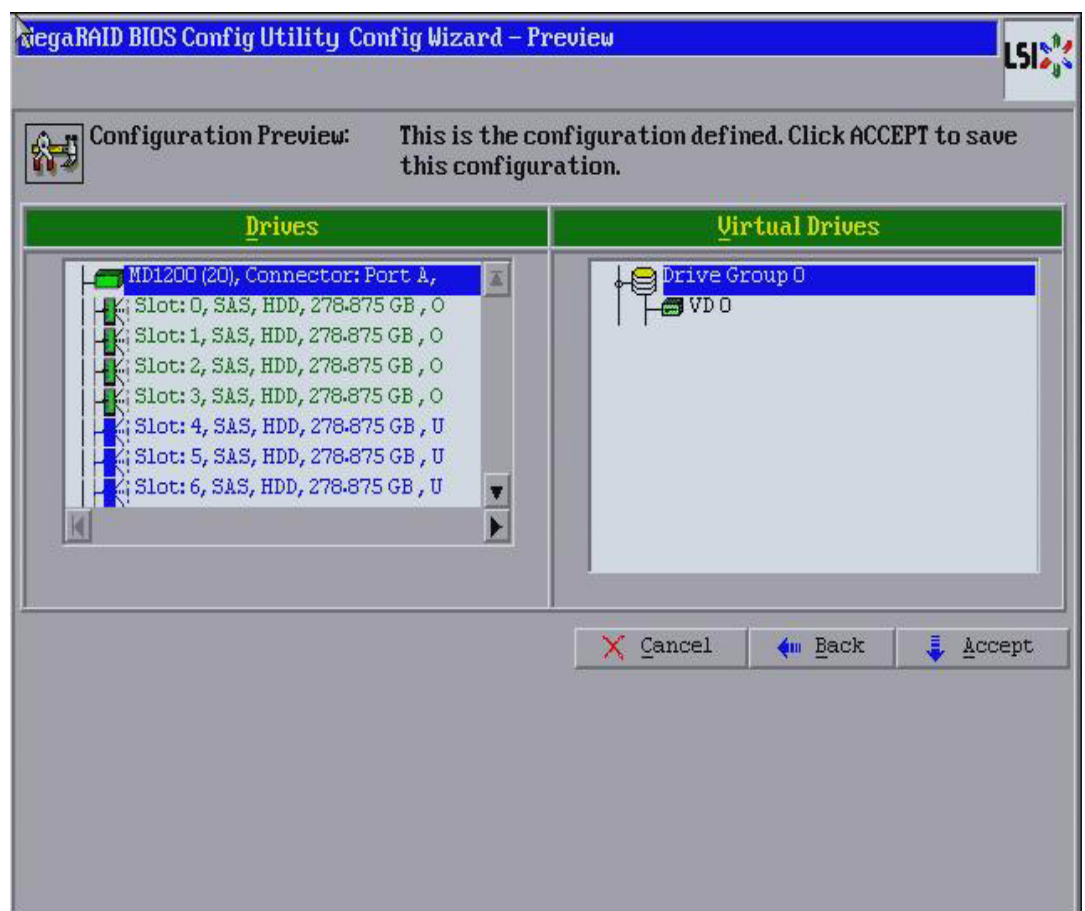


Figure 40: RAID 5 Configuration Preview Dialog

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

The WebBIOS main menu appears.

4.5.2.4 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 any 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.

NOTE: Integrated MegaRAID displays new drives as Just a Bunch of Disks (JBOD). For MegaRAID, unless the inserted drive contains valid DDF metadata, new drives display as JBOD. Rebuilds start only on Unconfigured Good drives, so you have to change the new drive state from JBOD to Unconfigured Good to start a rebuild.

When you select **Manual Configuration**, and click **Next**, the WebBIOS Drive Group Definition dialog appears. You use this dialog 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 41](#).
3. If you need to undo the changes, click **Reclaim**.
4. Choose whether to use power save mode.
5. Choose whether to use drive encryption.

The drop-down list in the **Encryption** field lists the options.

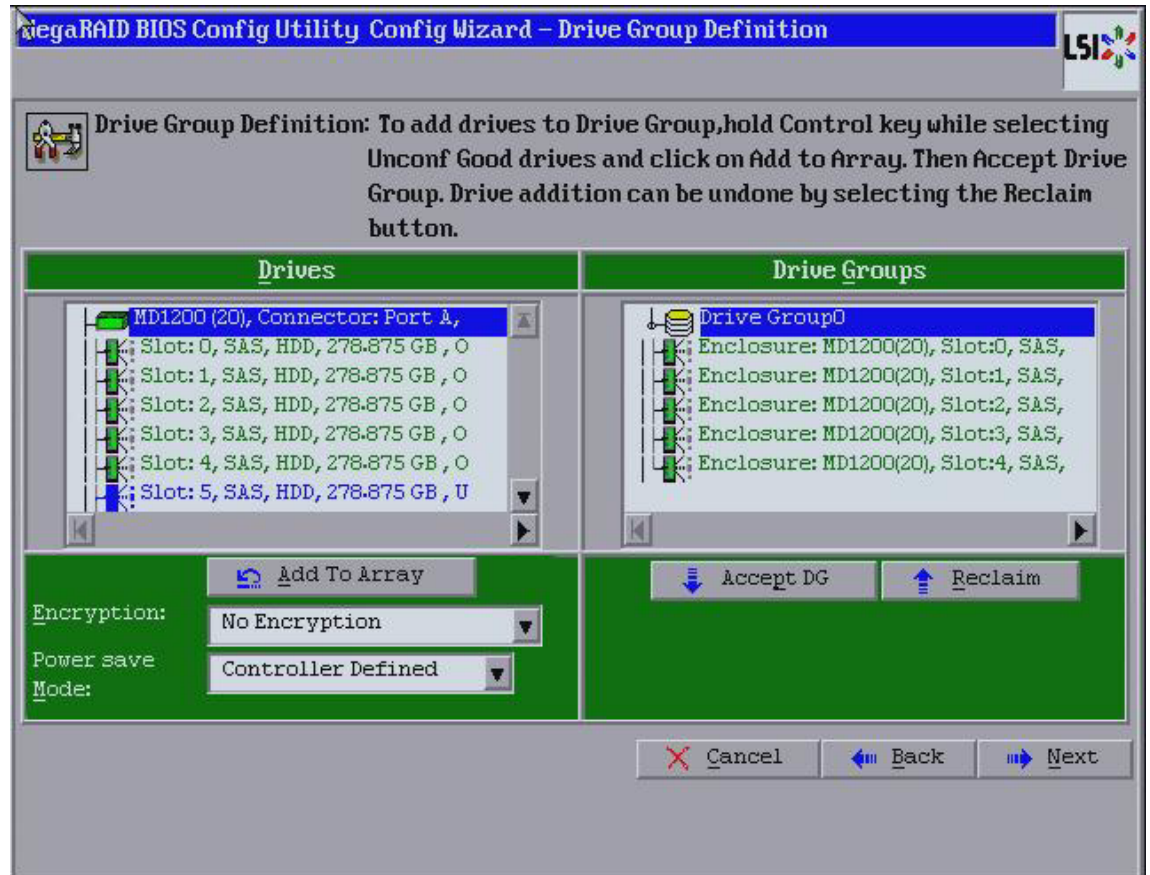


Figure 41: Drive Group Definition Dialog

6. After you finish selecting drives for the drive group, click **Accept DG** for each drive.
7. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 42](#). Use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

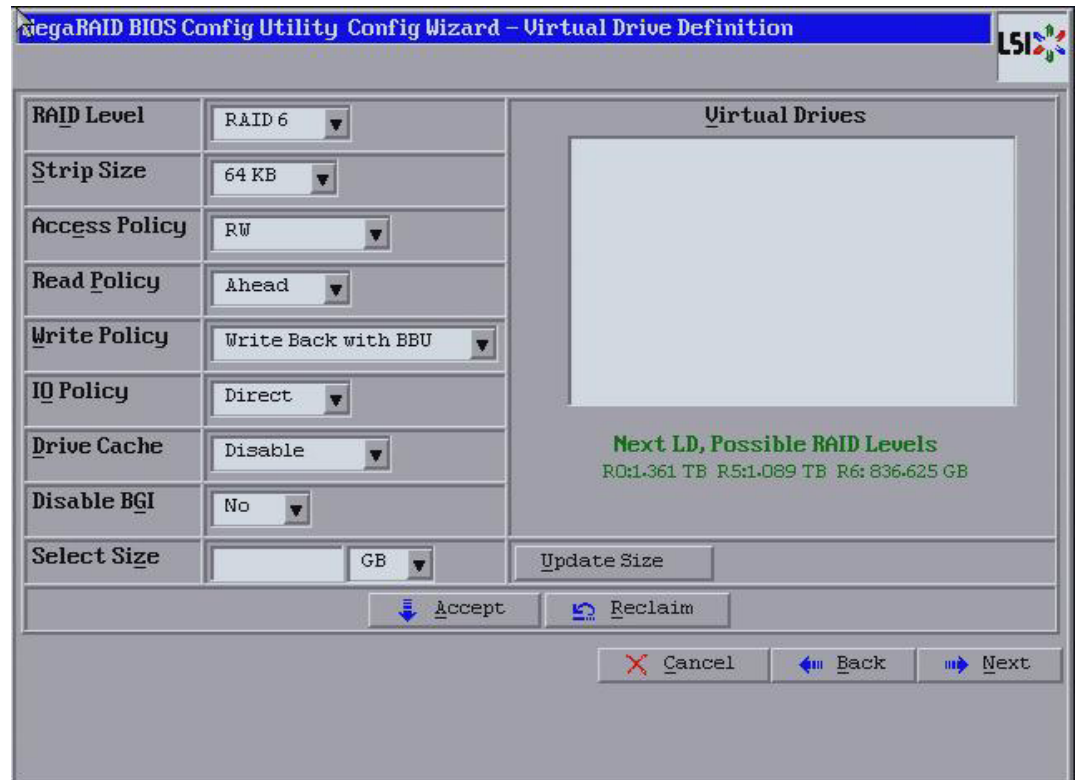


Figure 42: WebBIOS Virtual Drive Definition Dialog

8. Change the virtual drive options from the defaults listed on the dialog 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**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default setting is **64 KB**.

NOTE: WebBIOS does not allow you to select 8KB 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 option is the default setting.
 - Read Only:** Allow read-only access.
 - Blocked:** Do not allow access.

- **Read Policy:** Specify the read policy for this virtual drive.
 - Normal:** This option disables the read ahead capability. This option is the default setting.
 - 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This is the default.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. LSI recommends that you use *either* a battery to protect the controller cache, or a UPS to protect the entire system. If you do not use a battery or a UPS, and a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status:
 - No:** Leave background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This option is the default setting.
 - Yes:** Select **Yes** if you do not want to allow background initializations for configurations on this controller.

NOTE: New RAID 6 virtual drives require at least seven drives for a background initialization to start.

- **Select Size:** Specify the size of the virtual drive in MB, GB, or TB. Usually, this would be the full size for RAID 6 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.
 - **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels
9. Click **Accept** to accept the changes to the virtual drive definition.
If you need to undo the changes, click **Reclaim**.
 10. Click **Next** after you finish defining the virtual drives.

The Configuration Preview dialog appears, as shown in [Figure 43](#).

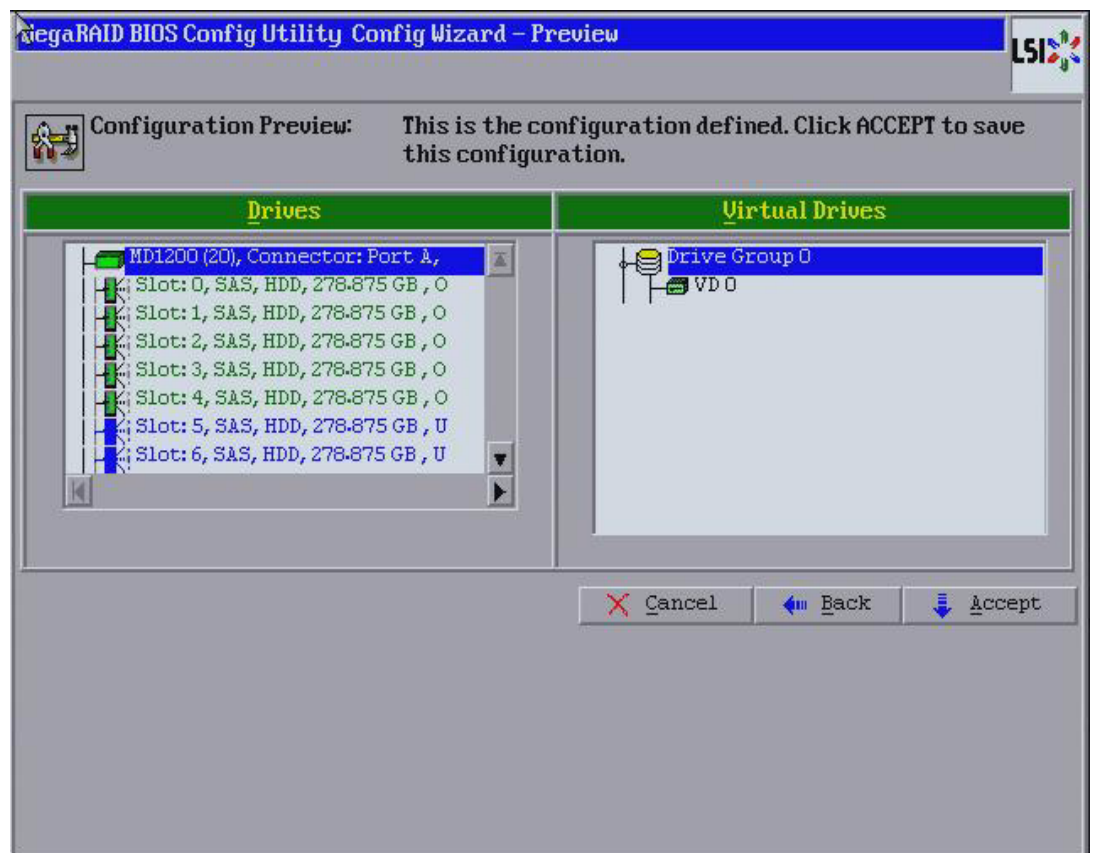


Figure 43: RAID 6 Configuration Preview Dialog

11. Check the information in the configuration preview dialog.
12. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
13. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.5.2.5 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 KB.

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 WebBIOS Drive Group Definition dialog appears.

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

1. Hold **Ctrl** key 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 **Drive Groups** panel on the right.
3. If you need to undo the changes, click **Reclaim**.
4. Click **Accept DG** to create a RAID 0 drive group.

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

5. Hold the **Ctrl** key while you select more ready drives in the Drives panel to create a second RAID 0 drive group.
6. Click **Add To Array** to move the drives to a second drive group configuration in the **Drive Groups** panel, as shown in [Figure 44](#).

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

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.)

7. Choose whether to use drive encryption.
8. Click **Accept DG** to create a RAID 0 drive group.

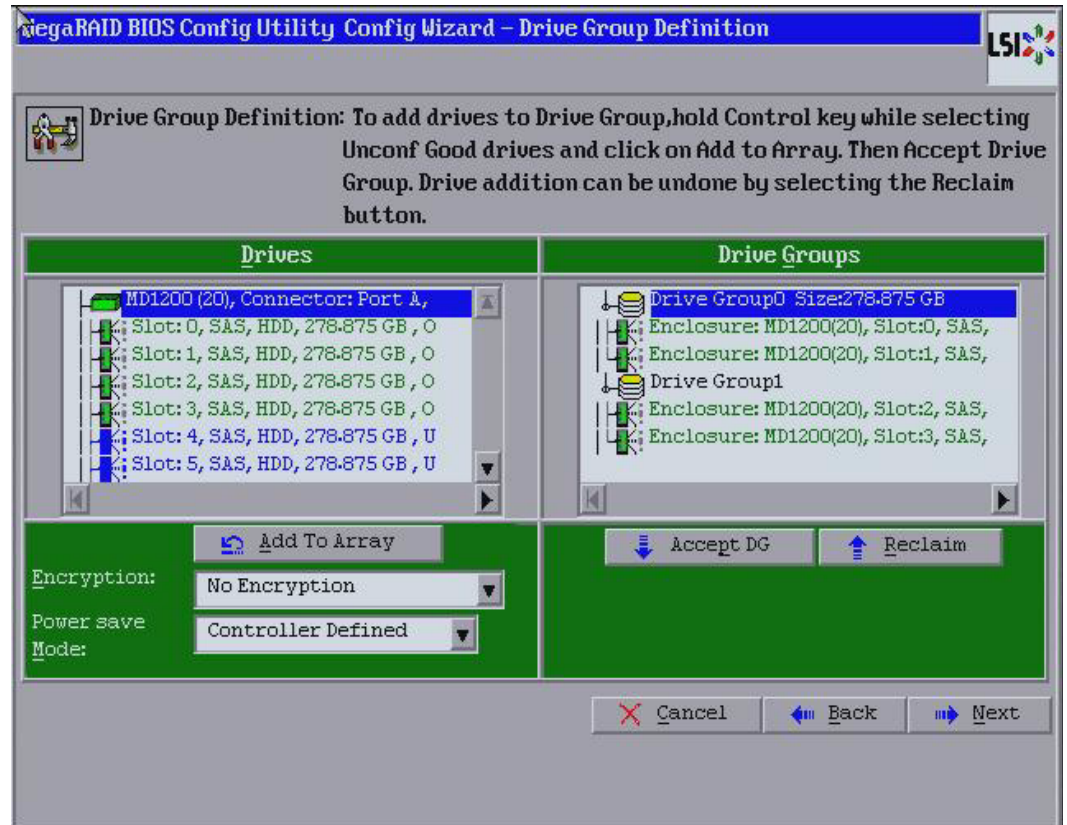


Figure 44: Drive Group Definition Dialog

9. Repeat step 4 through step 6 until you have selected all the drives you want for the drive groups.
10. After you finish selecting drives for the drive groups, select each drive group, and click **Accept DG** for each selection.
11. Click **Next**.

The Span Definition dialog appears, as shown in [Figure 45](#). This dialog shows the drive group holes that you can select to add to a span.

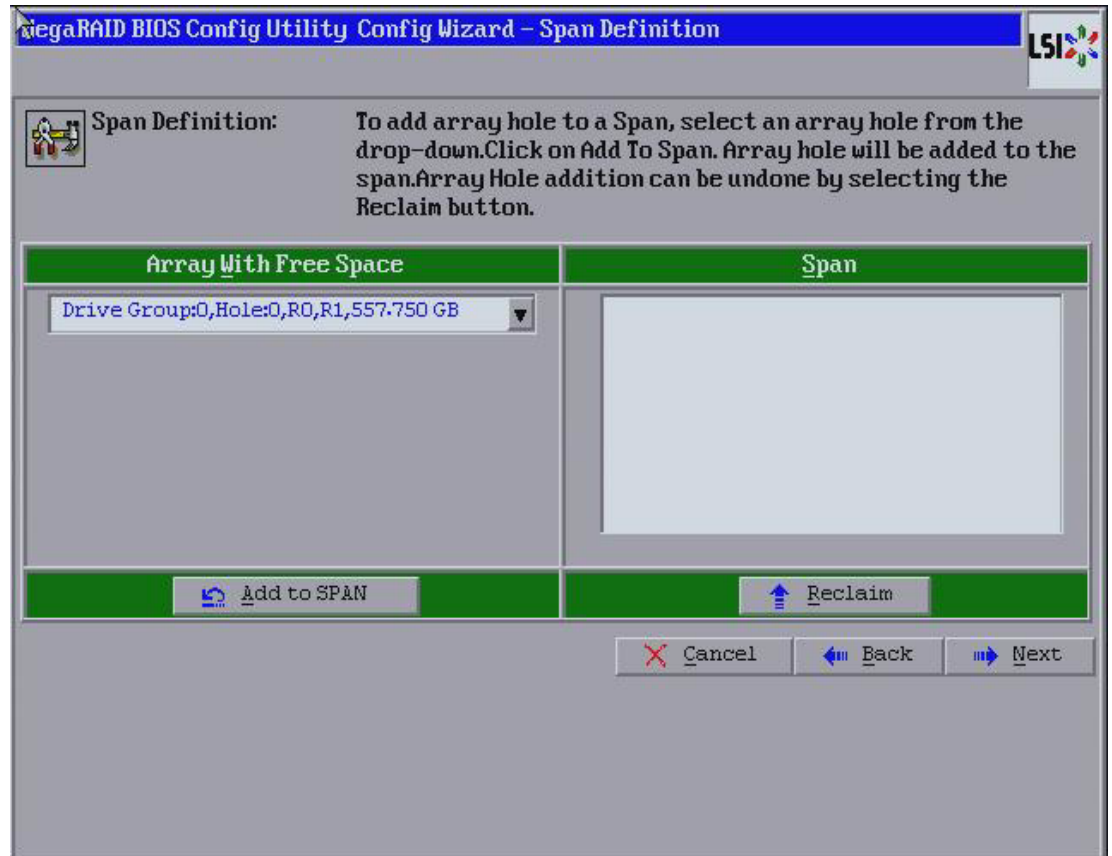


Figure 45: Span Definition Dialog

12. Under the Array With Free Space frame, select a drive group, and then click **Add to SPAN**.

The drive group you select appears in the right frame under Span.

13. Click **Add to SPAN**.
14. Repeat the previous two steps until you have selected all of the drive groups that you want.
15. Click **Next**.

The Virtual Drive Group Definition dialog appears, as shown in [Figure 46](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

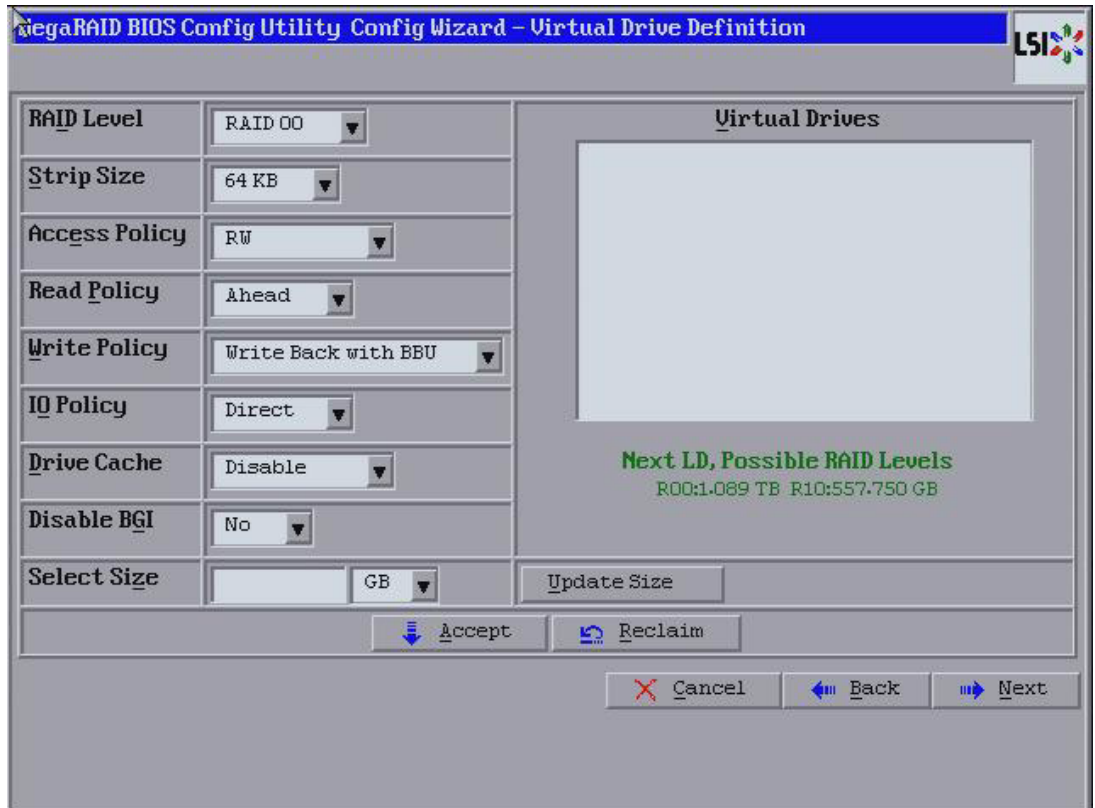


Figure 46: Virtual Drive Group Definition dialog

16. Change the virtual drive options from the defaults listed on the dialog as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down list shows the possible RAID levels for the virtual drive. Select **RAID 00**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default value is **64 KB**.
- **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 option 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 option 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This option is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Writeback mode to be used with or without a BBU. To protect the entire system, LSI recommends that you use *either* a battery to protect the controller cache or a UPS. 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) 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 setting 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 MB, GB, or TB. Usually, this would be the full size for RAID 00 shown in the **Configuration Panel** on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels
- 17. Click **Accept** to accept the changes to the virtual drive definition.
- 18. If you need to undo the changes, click **Reclaim**.
- 19. After you finish defining the virtual drives, click **Next**.

The Configuration Preview dialog appears, as shown in [Figure 47](#).

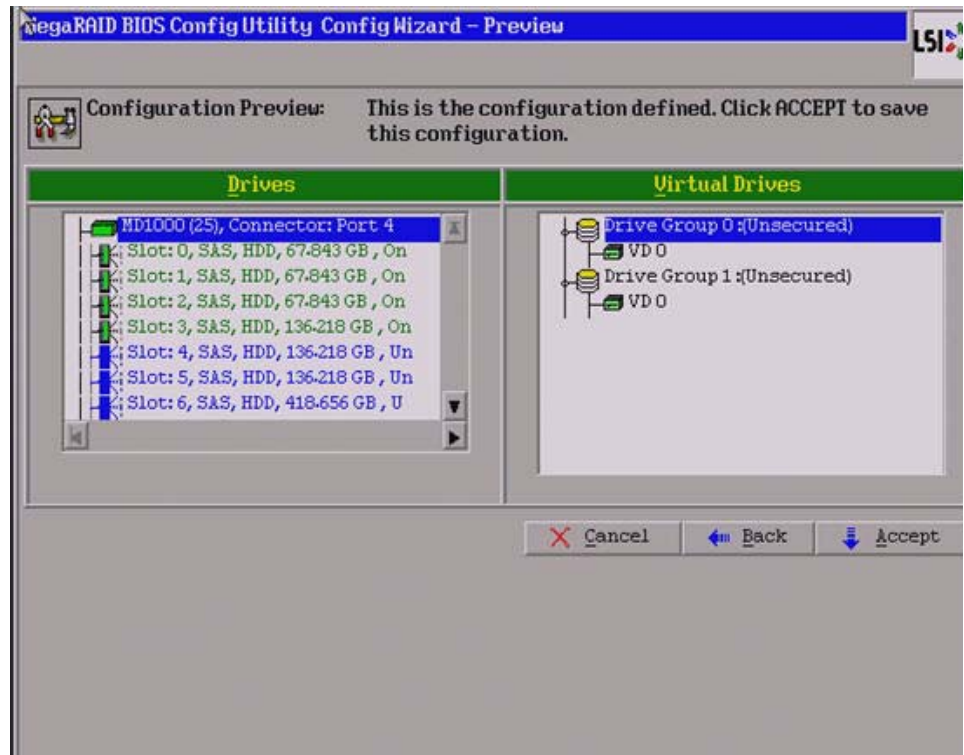


Figure 47: RAID 00 Configuration Preview Dialog

- 20. Check the information in the Configuration Preview Dialog.
- 21. 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 dialogs and change the configuration.
- 22. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.5.2.6 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 KB. 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 dialog appears.

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

1. Hold the **Ctrl** key 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 group configuration in the **Drive Groups** panel on the right.
3. If you need to undo the changes, click **Reclaim**.
4. Click **Accept DG** to create a RAID 1 drive group.

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

5. Click the icon for the next drive group to select it.
6. Hold the **Ctrl** key while selecting two more ready drives in the Drives panel to create a second RAID 1 drive group with two drives.
7. Click **Add To Array** to move the drives to a second two-drive group configuration in the Drive Groups panel, as shown in [Figure 48](#).

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

8. Choose whether to use power saving.
9. Choose whether to use drive encryption.

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.

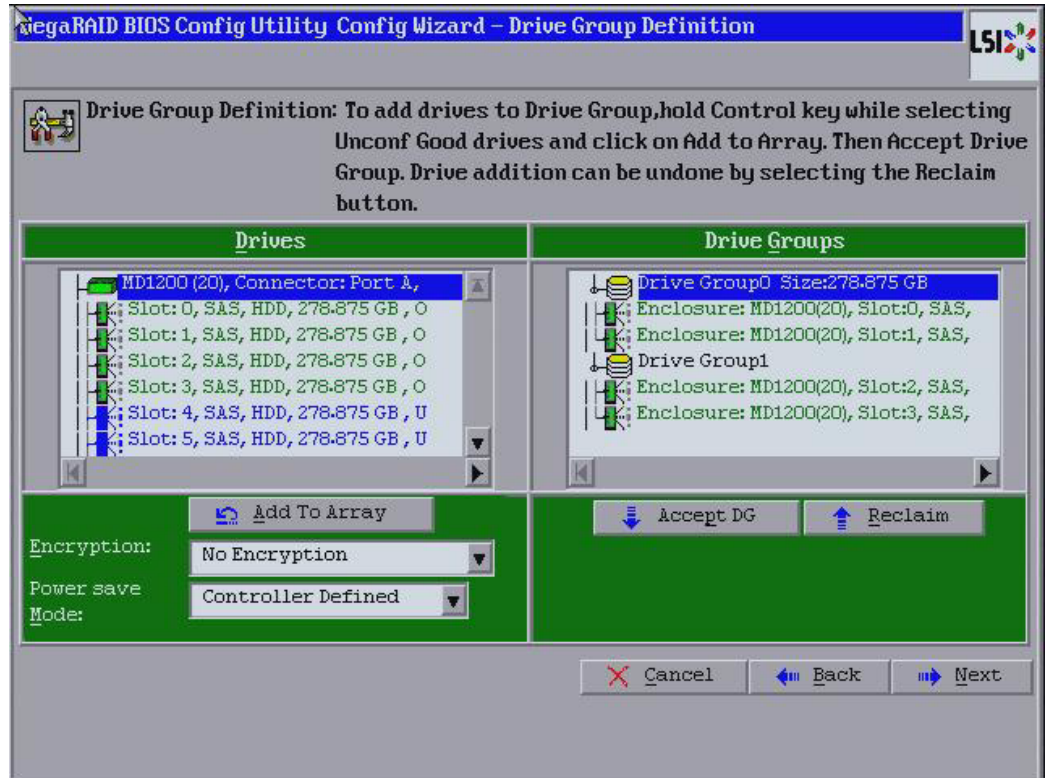


Figure 48: Drive Group Definition dialog

10. Repeat steps 7, 8, and 9 until you have selected all the drives you want for the drive groups.
11. After you finish selecting drives for the drive groups, select each drive group, and click **Accept DG** for each drive group.
12. Click **Next**.

The Span Definition dialog appears, as shown in [Figure 49](#). This dialog displays the drive group holes you can select to add to a span.

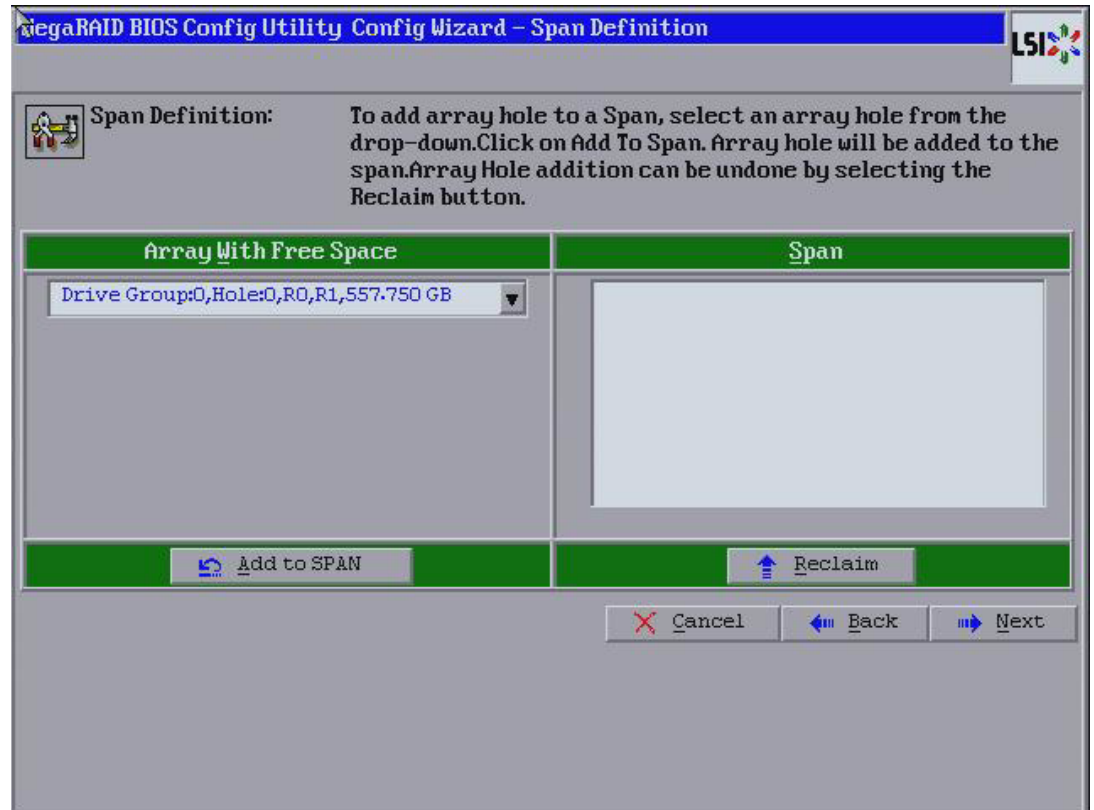


Figure 49: Span Definition Dialog

13. Under the **Array With Free Space** column, select a drive and click **Add to SPAN**.

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

14. Click **Add to SPAN**.

Both drive groups display in the right frame under Span.

15. If there are additional drive groups with two drives each, you can add them to the virtual drive.

16. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 50](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

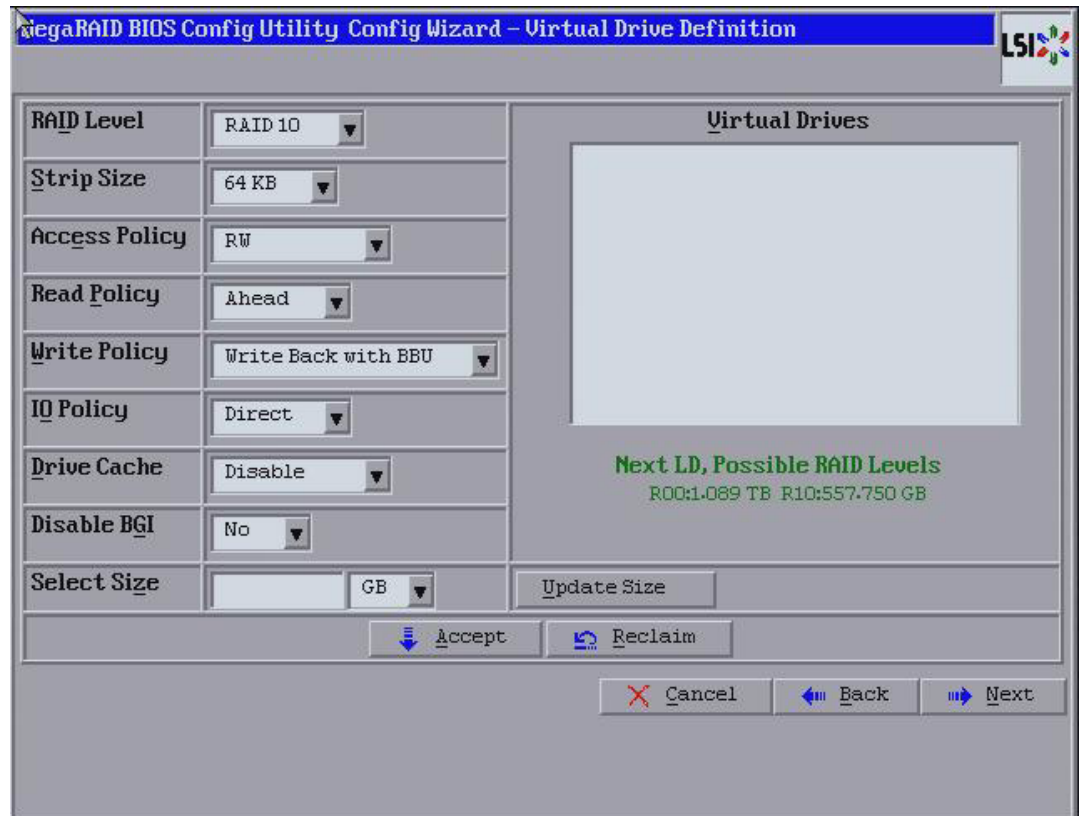


Figure 50: WebBIOS Virtual Drive Definition Dialog

NOTE: The WebBIOS Configuration Utility shows 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 2.

17. Change the virtual drive options from the defaults listed on the dialog 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**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is **64 KB**.

- **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 option is the default setting.
 - Blocked:** Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive.
 - Normal:** This option disables the read ahead capability. This option is the default setting.
 - 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This is the default.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. LSI recommends that you use *either* a battery to protect the controller cache, or a UPS to protect the entire system. If you do not use a battery or a UPS, and a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status.
 - No:** Leave background initialization enabled. This option means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This option is the default setting.
 - 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 MB, GB, or TB. Usually, this would be the full size for RAID 10 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.
 - **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels.
18. Click **Accept** to accept the changes to the virtual drive definition.
If you need to undo the changes, click **Reclaim**.
 19. After you finish defining the virtual drives, click **Next**.

The Configuration Preview dialog appears, as shown in [Figure 51](#).

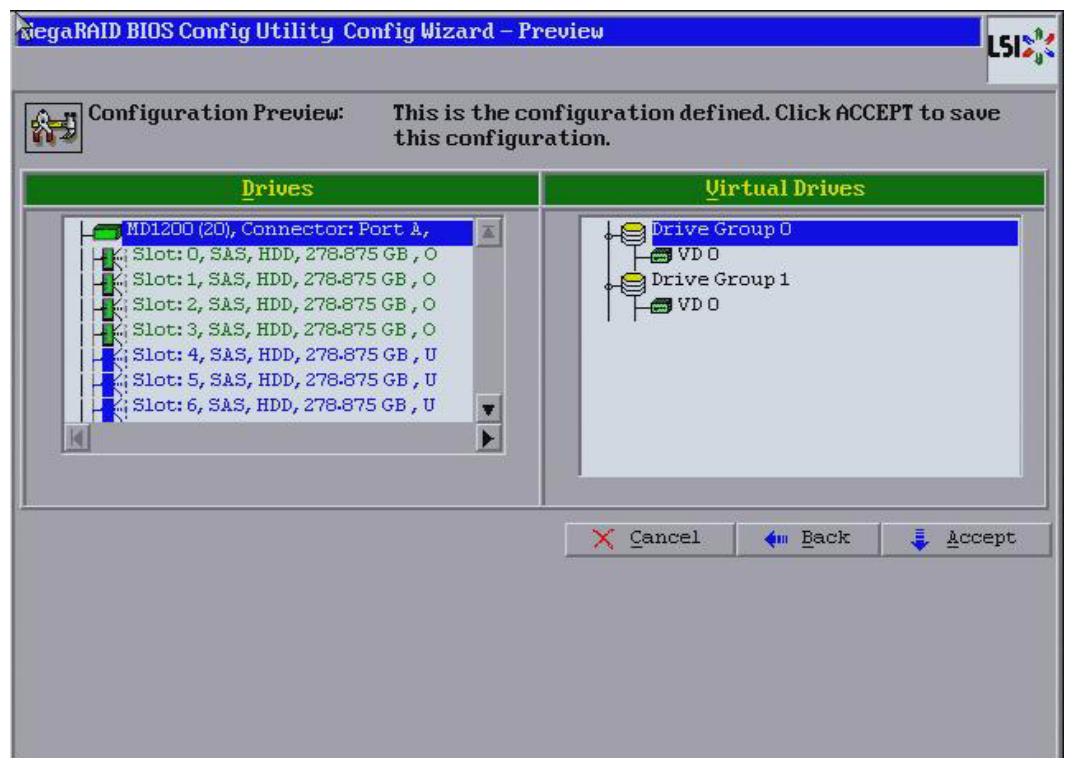


Figure 51: RAID 10 Configuration Preview Dialog

20. Check the information in the Configuration Preview.
21. 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 dialogs and change the configuration.
22. If you accept the configuration, click **Yes** at the prompt to save the configuration.
The WebBIOS main menu appears.

4.5.2.7 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 dialog appears. You use this dialog to select drives to create drive group.

1. Hold the **Ctrl** key 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 **Reclaim**.

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

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

4. Click the icon for the second drive group to select it.
5. Hold the **Ctrl** key 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 52](#).

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

7. Choose whether to use drive encryption.

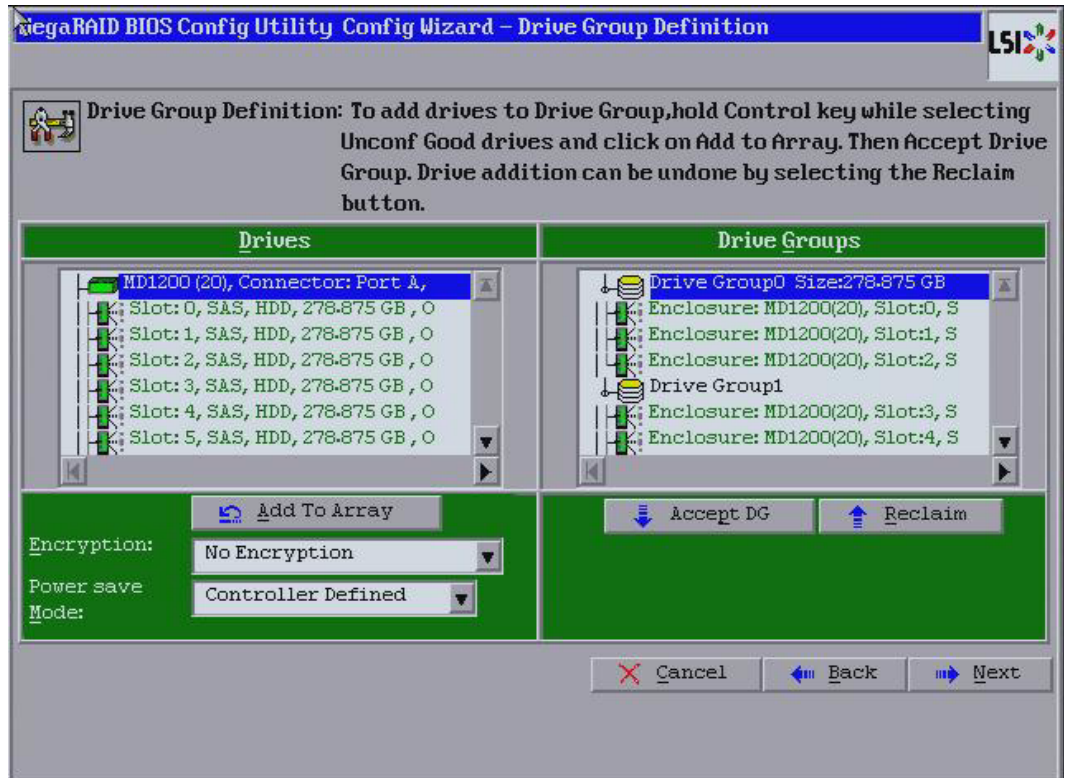


Figure 52: Drive Group Definition Dialog

8. After you finish selecting drives for the drive groups, select each drive group and click **Accept DG** for each drive group.
9. Click **Next**.

The Span Definition dialog appears, as shown in [Figure 53](#). This dialog displays the drive group holes you can select to add to a span.




Figure 53: Span Definition Dialog

10. Under the **Array With Free Space** column, 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.

11. Click **Add to SPAN**.

Both drive groups display in the right frame under Span.

12. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 54](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

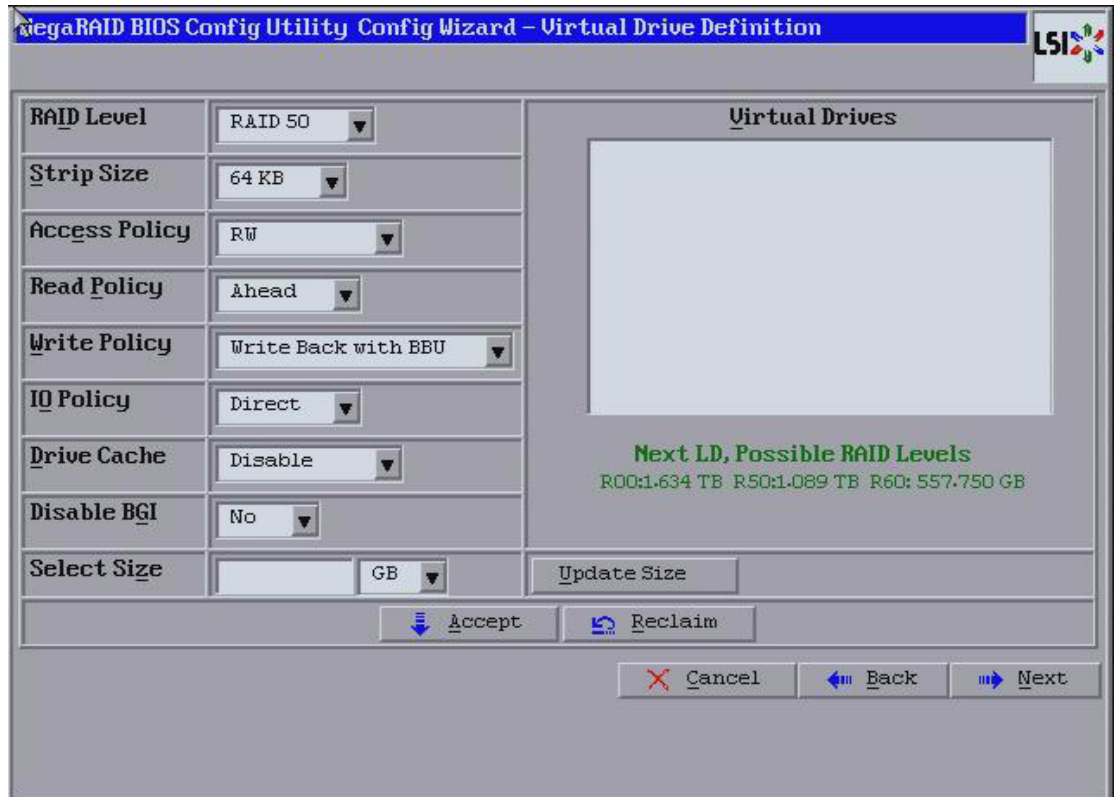


Figure 54: WebBIOS Virtual Drive Definition Dialog

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

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu list shows the possible RAID levels for the virtual drive. Select **RAID 50**.
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default setting is **64 KB**.
- **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 option is the default setting.
 - 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 option speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This setting is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. 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 a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) 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 option is the default setting.
 - 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 MB, GB, or TB. Usually, this option is the full size for RAID 50 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels
- 14. Either click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to undo the changes
- 15. Click **Next** after you finish defining the virtual drives.

The Configuration Preview dialog appears, as shown in [Figure 55](#).

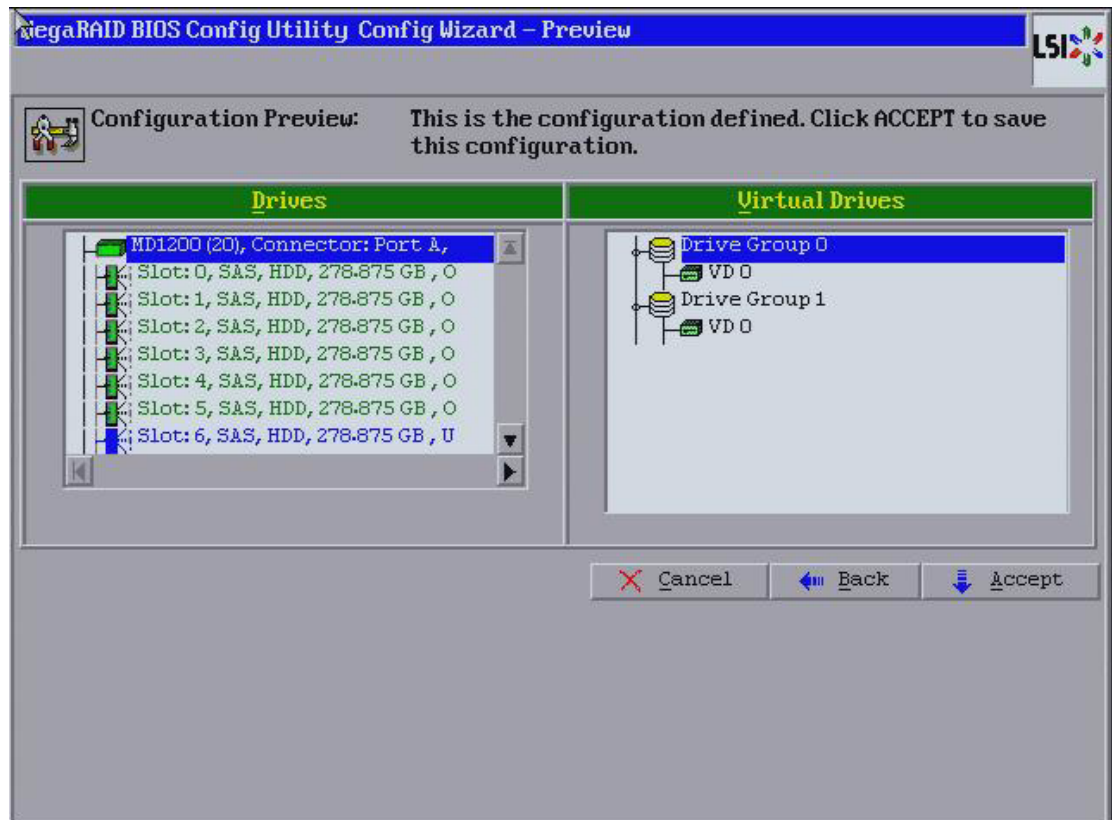


Figure 55: RAID 50 Configuration Preview Dialog

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

The WebBIOS main menu appears.

4.5.2.8 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 Drive Group Definition dialog appears. You use this dialog to select drives to create drive groups.

1. Hold the **Ctrl** key 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 **Reclaim**.

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

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

4. Click on the icon for the second drive group to select it.
5. Hold the **Ctrl** key 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 56](#).
7. If you need to undo the changes, click **Reclaim**.

Choose whether to use power saving.

8. Choose whether to use drive encryption.

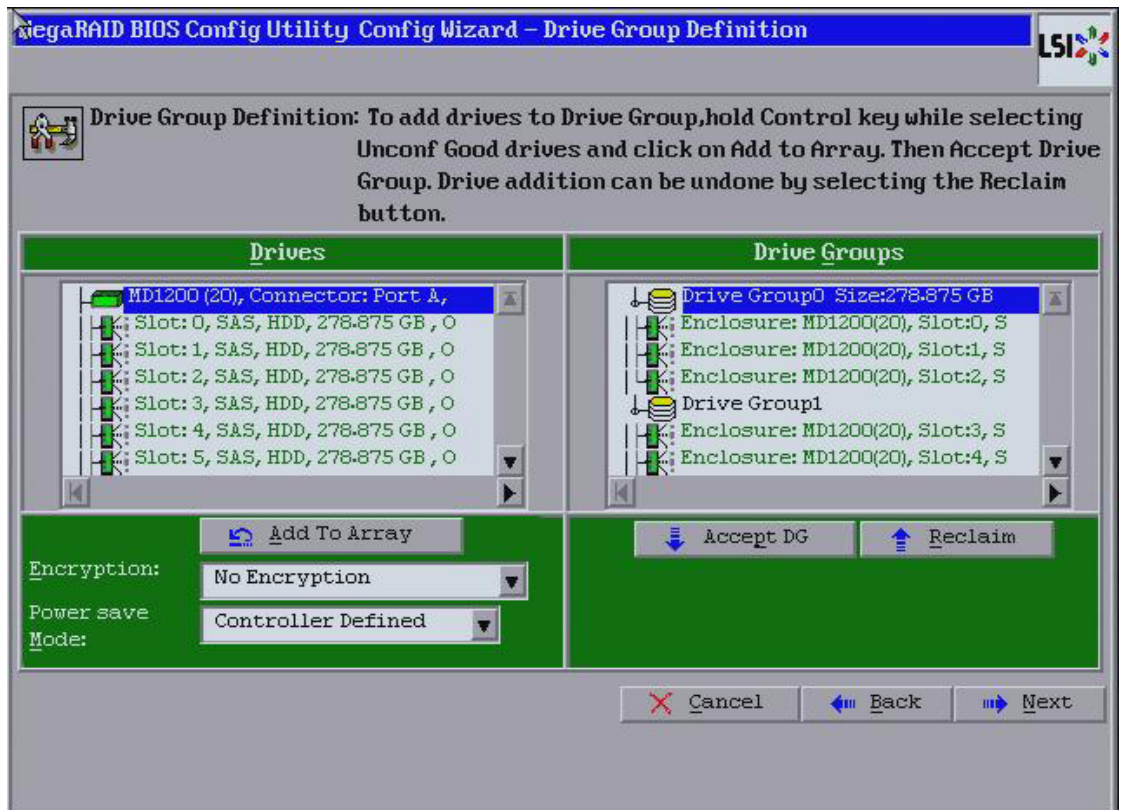


Figure 56: Drive Group Definition Dialog

9. After you finish selecting drives for the drive groups, select each drive group and click **Accept DG** for each drive group.
10. Click **Next**.

The Span Definition dialog appears, as shown in [Figure 57](#). This dialog displays the drive group holes you can select to add to a span.

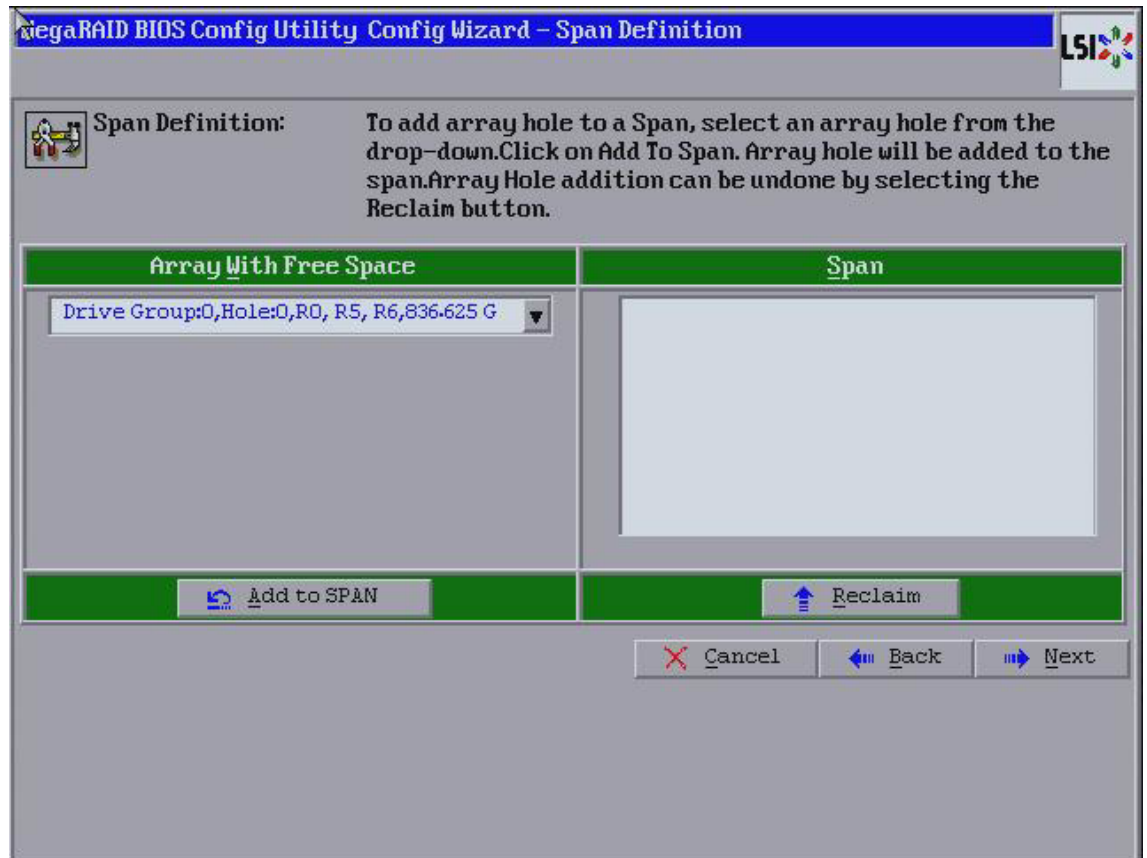


Figure 57: WebBIOS Span Definition Dialog

11. Under the heading **Array With Free Space**, select a drive group or drives, and click **Add to SPAN**.

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

12. Click **Add to SPAN**.

Both drive groups appear in the right frame under **Span**.

13. Click **Next**.

The Virtual Drive Definition dialog appears, as shown in [Figure 58](#). You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

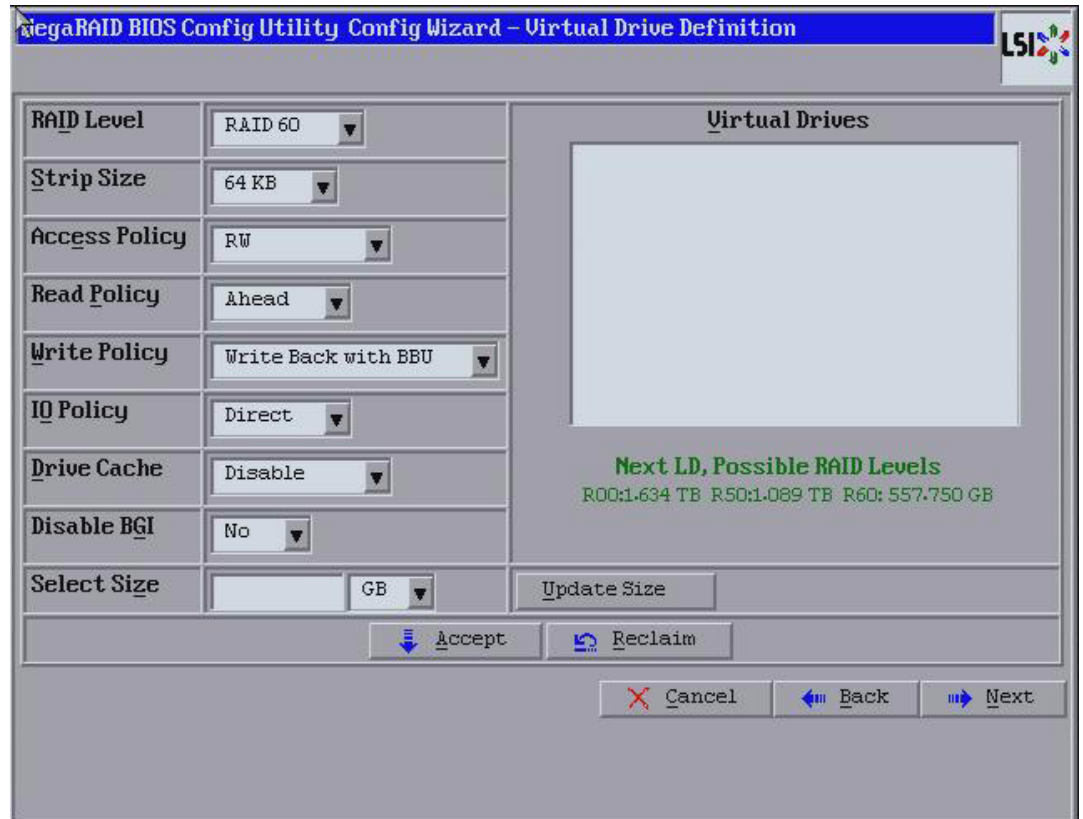


Figure 58: WebBIOS Virtual Drive Definition Dialog

14. Change the virtual drive options from the defaults listed on the dialog as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down list shows the possible RAID levels for the virtual drive. Select **RAID 60**.
- **Stripe Size:** The stripe size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists 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 strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, or 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is **64 KB**.
- **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 option is the default setting.
 - Blocked:** Do not allow access.

- **Read Policy:** Specify the read policy for this virtual drive.
 - Normal:** This setting disables the read ahead capability. This option is the default setting.
 - Ahead:** This setting 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 setting speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive.
 - WBack:** 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. This setting is recommended in Standard mode.
 - WThru:** 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. This option is the default setting.
 - Write Back with BBU:** Select this mode if you want the controller to use Write back 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 Write through mode if it detects a bad or missing BBU.

CAUTION: LSI allows Write back mode to be used with or without a BBU. 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 a power failure occurs, 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 option is the default setting.
 - 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 option is the default setting.
 - NoChange:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status.
 - No:** Leave background initialization enabled. This option means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This option is the default setting.
 - 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 MB, GB, or TB. Usually, this would be the full size for RAID 60 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels

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

15. Click **Accept** to accept the changes to the virtual drive definition.

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

16. Click **Next** after you finish defining virtual drives.

The webBIOS Config Utility Config Wizard Configuration Preview dialog appears, as shown in [Figure 59](#).

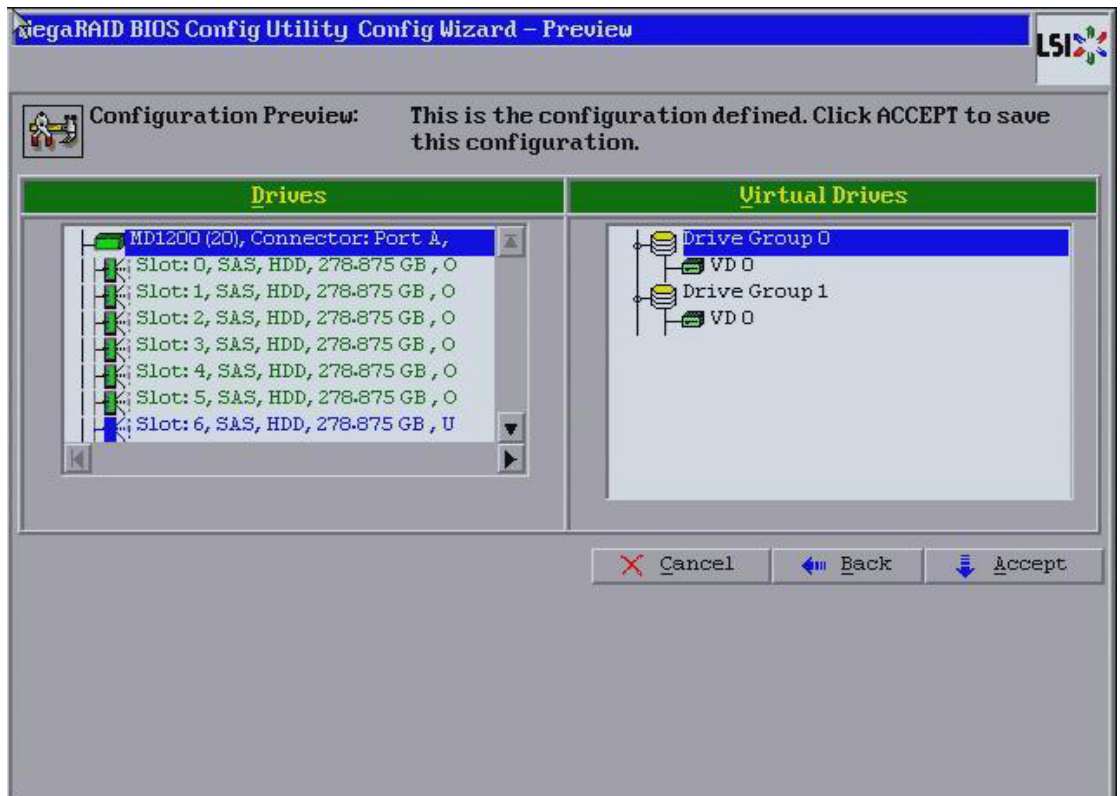


Figure 59: RAID 60 Configuration Preview Dialog

17. Check the information in the Configuration Preview dialog.

18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.

19. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.6 CacheCade Configuration

This section contains the procedures for creating CacheCadeRAID virtual drives for the CacheCade advanced software feature.

The MegaRAID CacheCade advanced software improves application performance by expanding the MegaRAID read caching capacity and introducing write caching. The CacheCade feature uses high-performing CacheCade software as a secondary tier of cache to provide faster reads and to maximize transactional I/O performance.

NOTE: This procedure does not create a RAID configuration. It creates a CacheCade software virtual drive that functions as a secondary tier of cache.

Using the CacheCade 2.0 SSD Read Caching software as controller cache allows for very large data sets to be present in cache, delivering performance up to 50 times greater than regular cache in read-intensive applications, such as online transaction processing (OLTP), and file and Web server workloads. The solution accelerates the IO performance of HDD-based drive groups while requiring only a small investment in CacheCade software technology.

To support full-throughput for multiple direct-attached CacheCade software, this feature reduces I/O-processing overhead in the 2208-chip-based MegaRAID controllers. CacheCade offers performance equivalent to flash-based controllers and better performance for RAID 5 and RAID 6 when compared to Fusion I/O.

4.6.1 Creating a CacheCade 2.0 SSD Read Caching Configuration

CacheCade 2.0 SSD Read Caching software provides you with read caching capability.

Perform the following steps to create a CacheCade 2.0 drive group:

1. Click **Configuration Wizard** on the WebBIOS main dialog.

The Select Configuration Wizard dialog appears, as shown in [Figure 60](#).

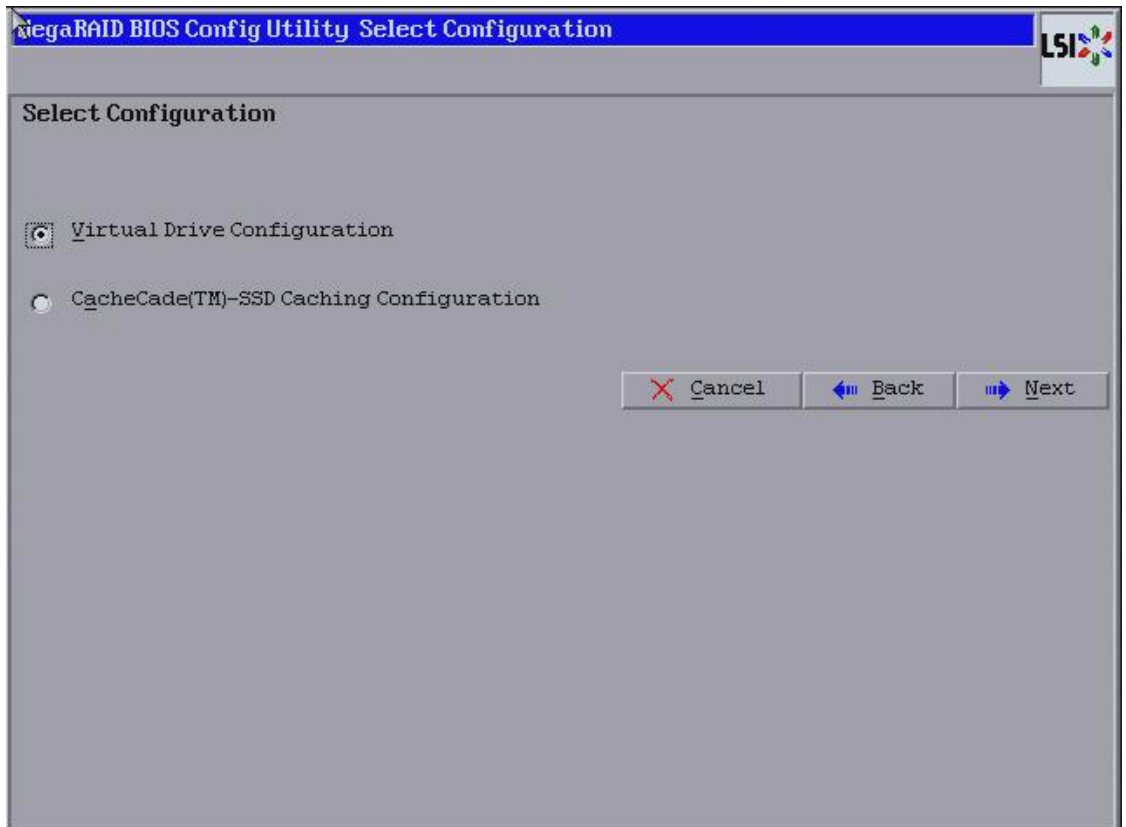


Figure 60: WebBIOS Select Configuration Wizard Dialog

2. Select the **CacheCade(TM)-SSD Caching Configuration** and click **Next**. The Drive Group Definition dialog appears.

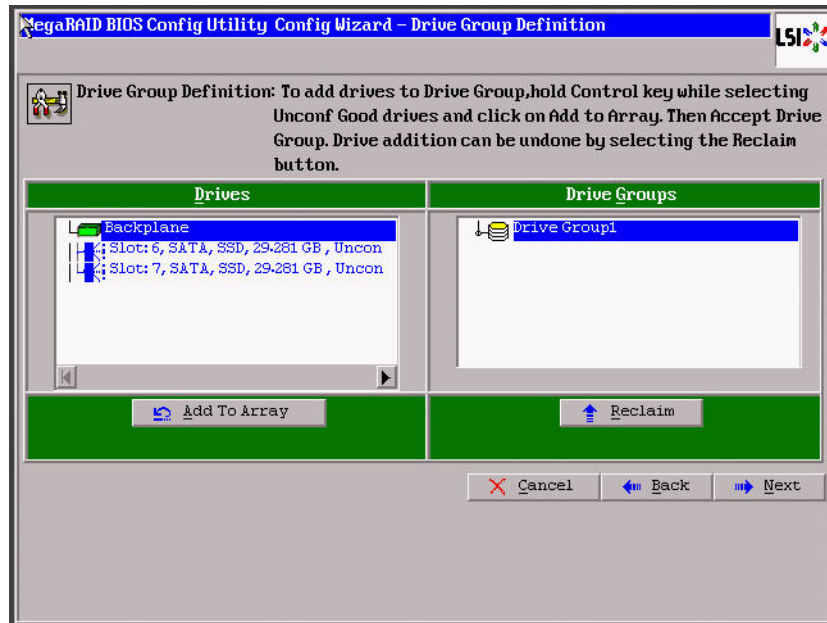


Figure 61: Drive Group Definition

3. Select a drive from the left frame, and click **Add To Array**. The selected drive now appears in **Drive Groups**, and the **Accept DG** button appears.
4. Click **Accept DG**. A drive group is created and appears in **Drive Groups**.
5. Click **Next**. The Span Definition dialog appears.

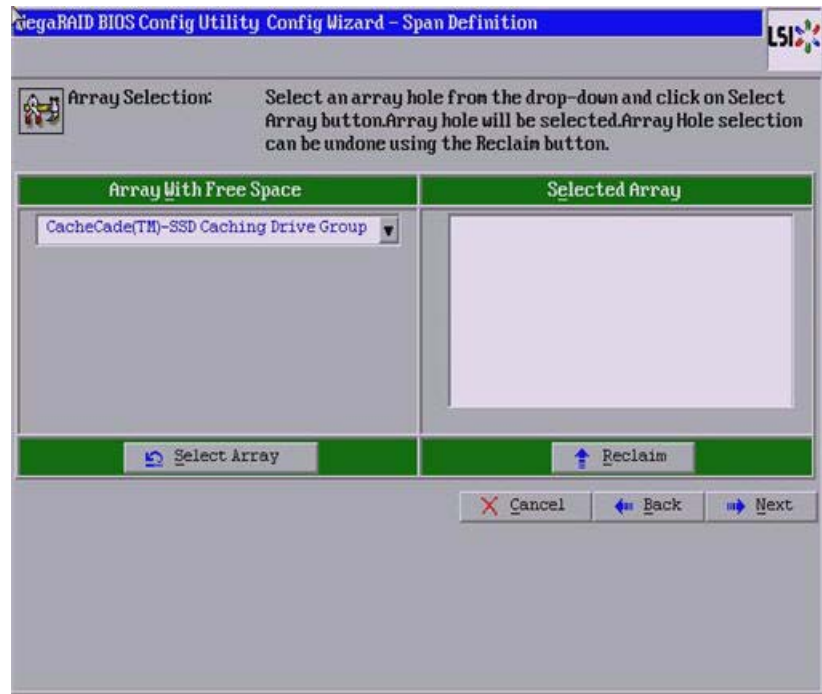


Figure 62: Span Definition

6. Select an array with free space from the drop-down list, and click **Select Array**.
The selected array moves to the right frame under the heading **Selected Array**.
7. Click **Next**.
The Create CacheCade - SSD Caching Disk dialog appears, as shown in [Figure 63](#).




Figure 63: SSD Caching Disk Dialog

8. Click **Accept** to accept the drive group.
If you need to undo the changes, click **Reclaim**.
The Config Wizard-Preview dialog appears, as shown in [Figure 64](#).

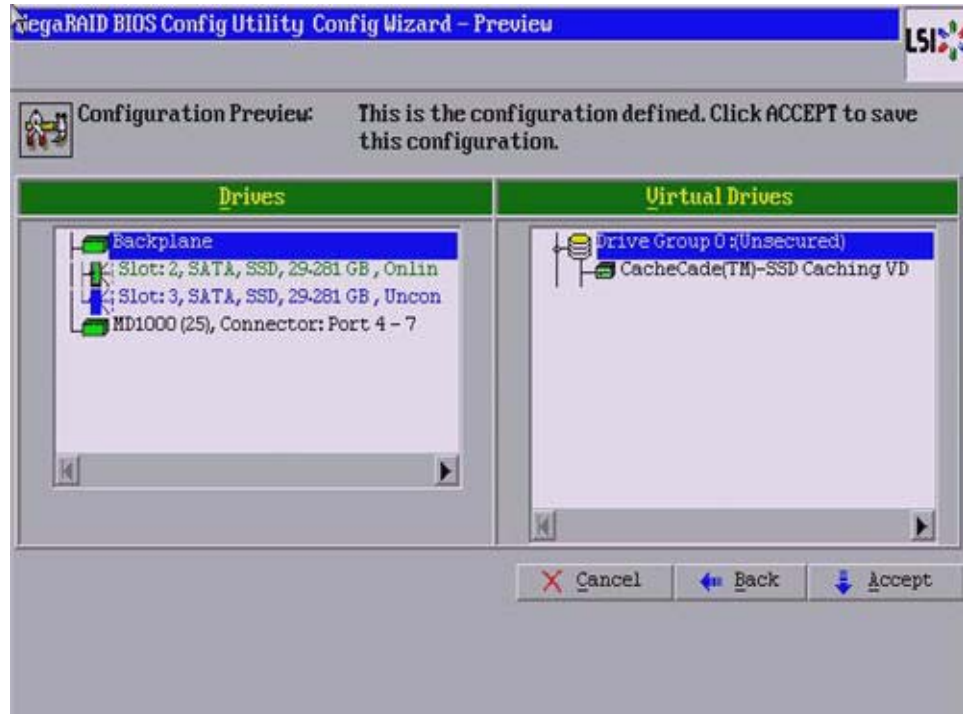


Figure 64: CacheCade Configuration Preview Dialog

9. Click **Accept** if the configuration is correct. Otherwise, or click **Back** to return to the previous dialogs and change the configuration.
10. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS Config Utility Virtual Configuration dialog, as shown in [Figure 65](#). It shows the CacheCade virtual drive.

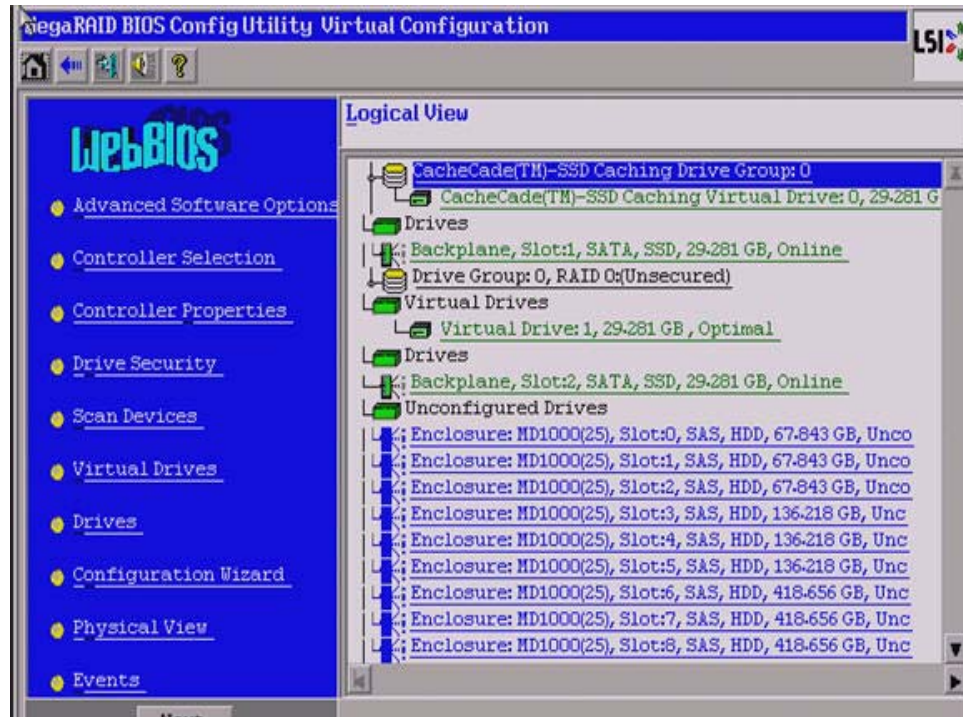


Figure 65: WebBIOS Main Menu with a CacheCade Virtual Drive

4.6.2 Creating a CacheCade Pro 2.0 SSD Read/Write Caching Configuration

CacheCade Pro 2.0 SSD Read/Write Caching software provides you with read and write caching capability.

Perform the following steps to create a CacheCade Pro 2.0 drive group:

1. Click **Configuration Wizard** on the WebBIOS CU main screen. The WebBIOS Configuration Wizard appears.

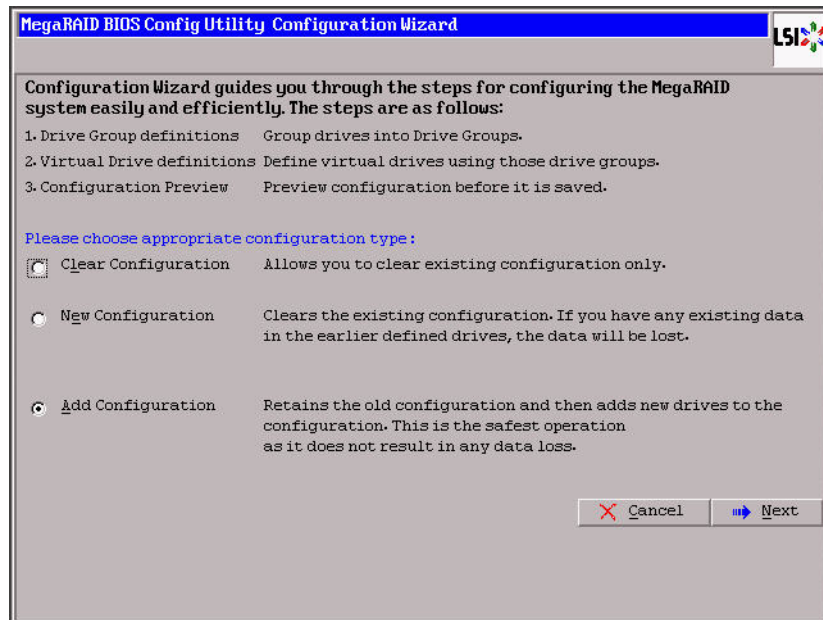


Figure 66: WebBIOS Configuration Wizard

2. **Add Configuration** is selected by default. Click **Next**. The Select Configuration screen appears.

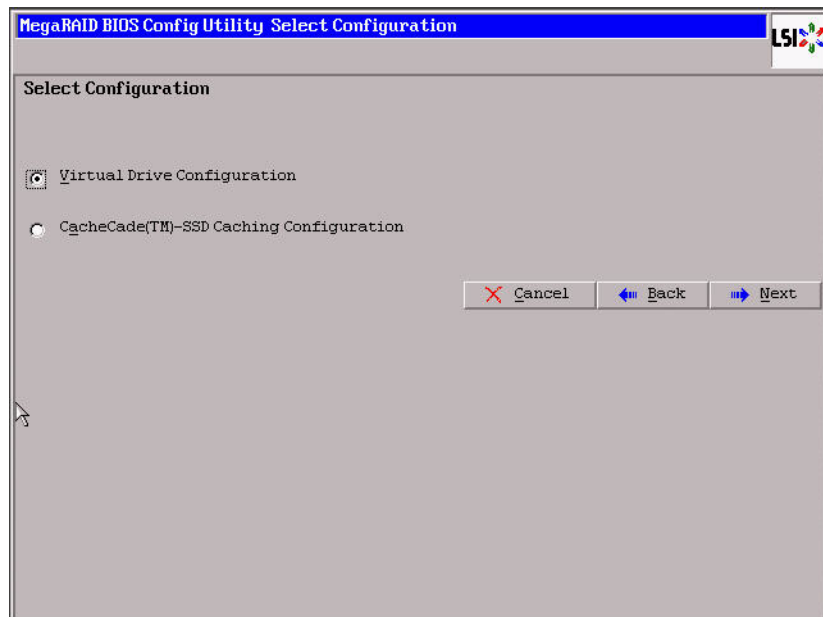


Figure 67: Select Configuration

3. Select **CacheCade(TM) - SSD Caching Configuration** and click **Next**.

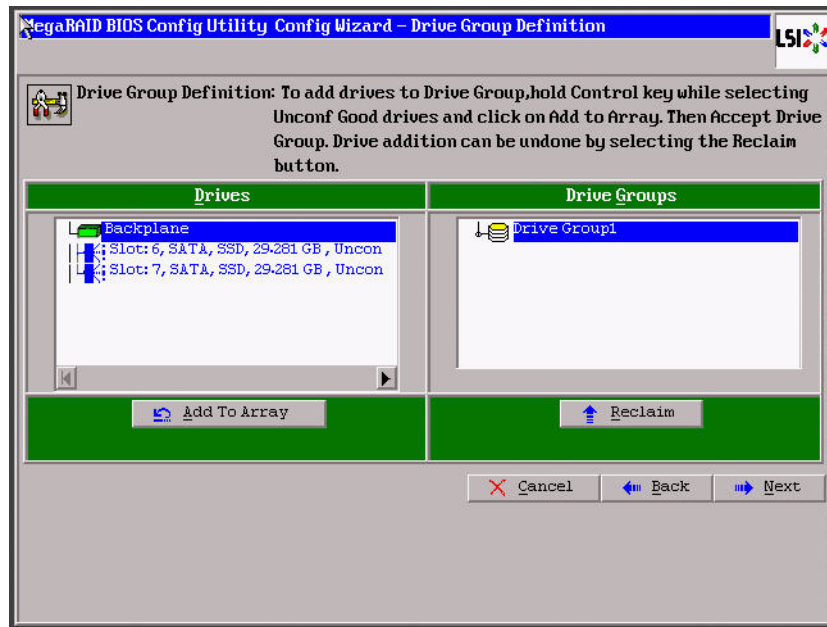


Figure 68: Drive Group Definition

4. Select a drive from the left frame, and click **Add To Array**. The selected drive now appears in **Drive Groups**, and the **Accept DG** button appears.
5. Click **Accept DG**. A drive group is created and appears in **Drive Groups**.
6. Click **Next**. The Span Definition screen appears.

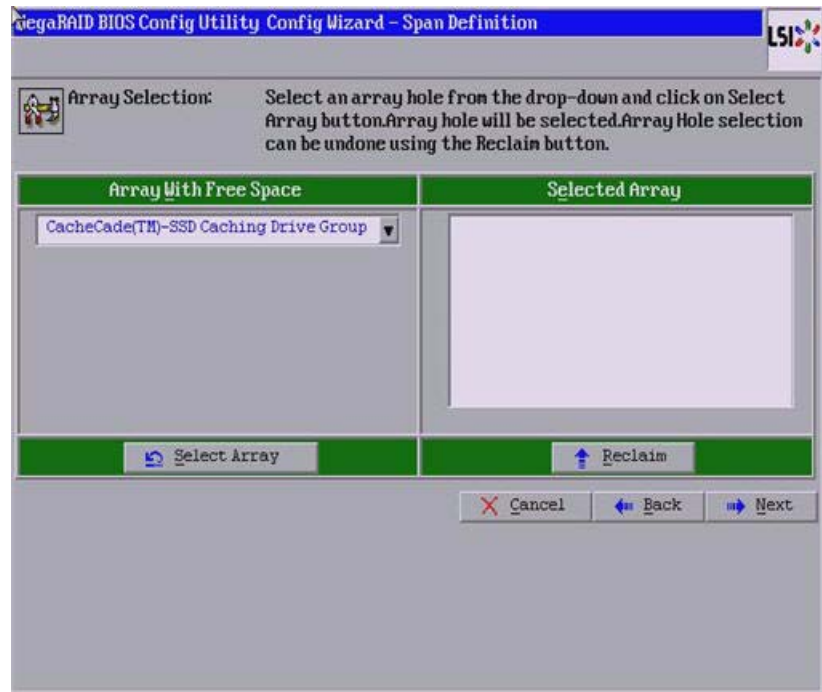


Figure 69: Span Definition

7. Select an array from **Array With Free Space** and click **Select Array**.
8. Click **Next**.
9. Select a RAID level from the **RAID Level** drop-down list.
10. Select a write policy from the **Write Policy** drop-down list.
 - Write Back: In Write Back mode, the CacheCade virtual drive is used for both read as well as write cache. However, if the CacheCade virtual drive becomes degraded, the CacheCade virtual drive will be used only as read cache. This is the default write policy.
 - Write Through: In Write Through mode, the CacheCade virtual drive is used as read only cache.
 - Always Write Back: In Always Write Back mode, the CacheCade virtual drive is used for both read and write cache.
11. Click **Accept**. A confirmation screen appears asking you to confirm your selections.
12. Click **Yes** to confirm and click **Next**. The Configuration Preview screen appears.
13. Click **Accept**. A confirmation screen appears asking if you want to save your configuration.
14. Click **Yes**. Your configuration is saved, and you are taken back to the WebBIOS CU main screen. The new CacheCade drive group appears in the frame under **Logical View**.

4.6.2.1 Modifying CacheCade Pro 2.0 Virtual Drive Properties

You can modify the default write policy of a CacheCade virtual drive. You can also delete a CacheCade virtual drive.

Perform the following steps to modify the CacheCade virtual drive properties:

1. In the WebBIOS CU main screen, click the CacheCade virtual drive whose properties you want to modify. The CacheCade Virtual Drive Properties screen appears.

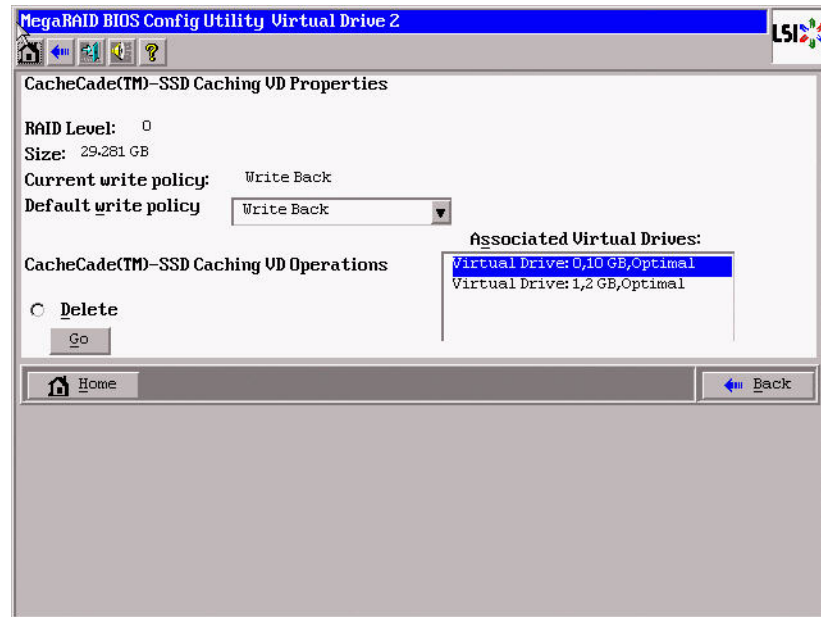


Figure 70: CacheCade Virtual Drive Properties

2. To modify the default write policy of a Cachecade virtual drive, select a write policy from the **Default write policy** drop down list.
3. To delete a CacheCade virtual drive, click **Delete** and then click **Go**.
4. Click **Home** to return to the WebBIOS CU main screen.

4.6.2.2 Enabling/Disabling SSD Caching on a Virtual Drive

You can associate a virtual drive to a cache pool or not by enabling or disabling SSD caching on that virtual drive.

When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade virtual drive.

When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade SSD Caching virtual drive is removed. This option is only available when the virtual drive's caching is currently enabled.

Perform the following steps to enable/disable SSD caching on a virtual drive:

1. In the WebBIOS CU main screen, click a virtual drive. The Virtual Drive Properties screen appears.

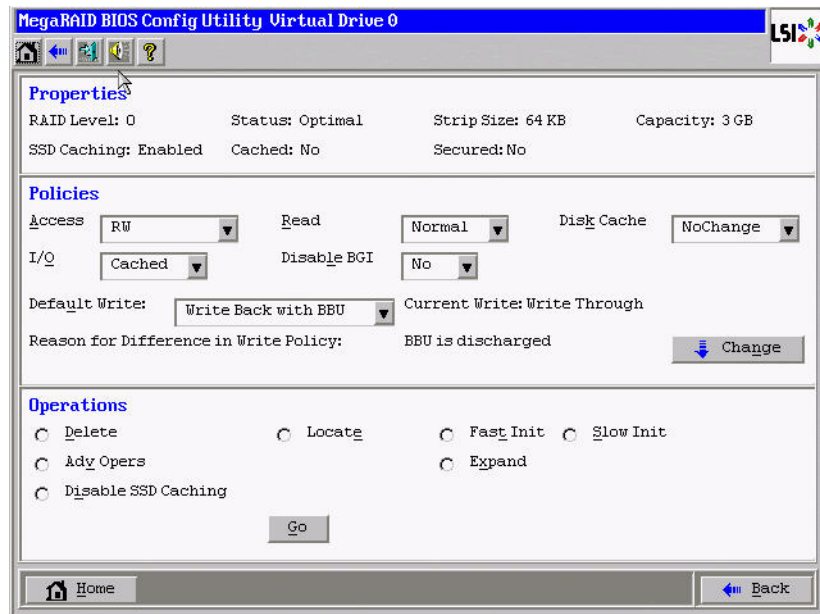


Figure 71: Virtual Drive Properties

2. Select the **Enable SSD Caching** or **Disable SSD Caching** radio button in the **Operations** frame.
3. Click **Go**. A confirmation page appears asking you to confirm your selection.
4. Click **Yes**.

4.6.2.3 Enabling/Disabling SSD Caching on Multiple Virtual Drives

You can enable and disable SSD caching on multiple virtual drives at one time.

When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade virtual drive. This option is only available when there is at least one virtual drive in the configuration.

When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade virtual drive is removed. This option is only available when there is at least one virtual drive in the configuration.

Perform the following steps to enable/disable SSD Caching on multiple virtual drives:

1. In the WebBIOS CU main screen, click **Controller Properties** in the left frame. The Controller Properties screen appears.
2. Click **Next** until you reach the last controller properties screen.
3. Click **Manage** (next to SSD Caching). The Manage SSD Caching screen appears.

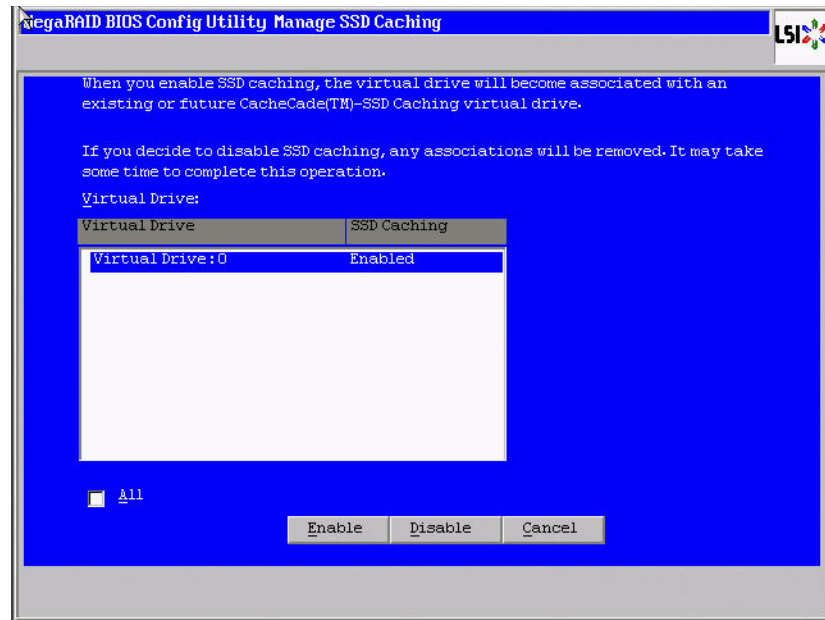


Figure 72: WebBIOS CU Manage SSD Caching

All virtual drives that have SSD caching enabled or disabled are listed.

4. Select the **All** check box and click **Enable** or **Disable**.
 - Click **Enable** to enable SSD caching on all the virtual drives that are currently disabled.
 - Click **Disable** to disable SSD caching on all virtual drives that are currently enabled.

A confirmation screen appears asking for your confirmation.

5. Click **Yes** to continue with disabling SSD caching on all virtual drives.

4.6.2.4 Enabling SSD Caching on New Virtual Drives

You can enable SSD caching on a virtual drive when the virtual drive is being created in the Create Virtual Drive wizard.

Once the virtual drive has been created using the wizard, the Manage SSD Caching screen appears:

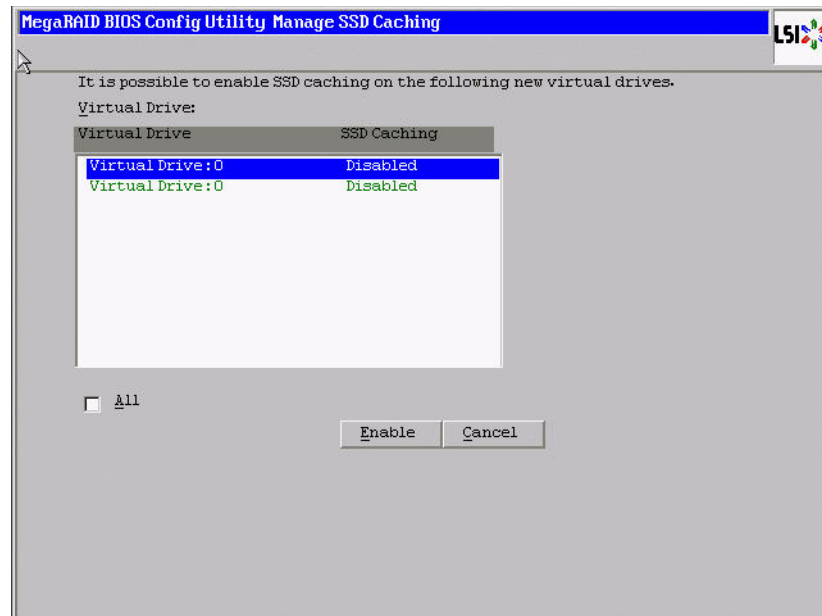


Figure 73: WebBIOS CU Manage SSD Caching

By default, all virtual drives are disabled.

Click **Enable** to enable SSD caching on a virtual drive. Select the **All** check box and click **Enable** to enable SSD caching on all virtual drives.

4.6.2.5 Clearing Configurations on CacheCade Pro 2.0 Virtual Drives

You can clear all the configurations on a CacheCade virtual drive.

1. In the WebBIOS CU main screen, click **Configuration Wizard**. The first screen of the configuration wizard appears.
2. Select **Clear Configuration**. A confirmation screen appears.

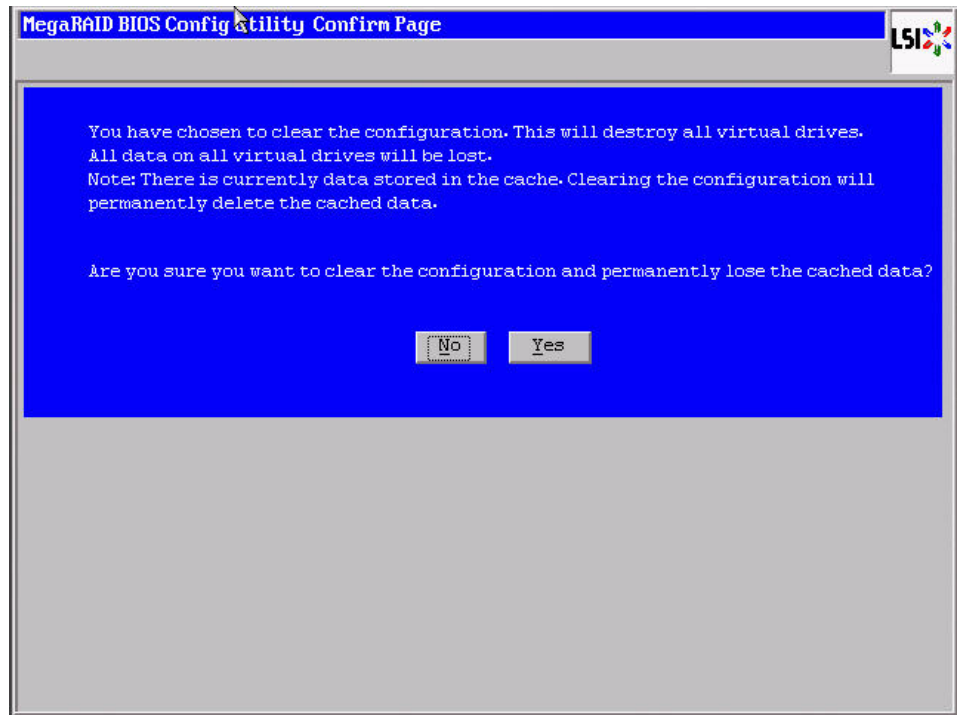


Figure 74: Confirmation Page

3. Click **Yes** to clear the configuration.

If the cache becomes inconsistent before the clear configuration operation is performed, the firmware returns an error code. The Confirm Loss of Cache dialog appears as a follow up dialog to the Confirm Clear Configuration dialog. Click **Yes** to clear the configuration.

4.6.2.6 Removing Blocked Access

At times, an error may occur in the CacheCade virtual drive and this causes a blocked access to the associated virtual drive.

It is advisable to wait for sometime for the error in the CacheCade virtual drive to get sorted. You can also try to solve the error in the CacheCade virtual drive and bring it back to an optimal status. Once the CacheCade virtual drive is in an optimal status, the blocked virtual drive returns to its former access policy automatically.

The text *Access Blocked* gets appended next to the *Optimal* status of the affected virtual drive in the WebBIOS CU main screen.

1. Click on the affected virtual drive to view the Virtual Drive Properties screen. The **Access** field displays *Blocked* as the access policy, as shown in the following figure.

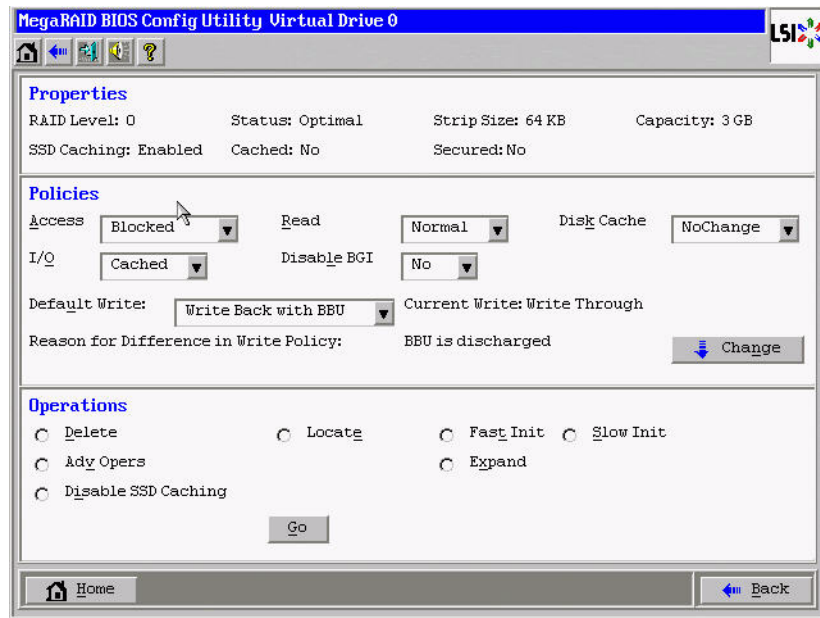


Figure 75: Virtual Drive Properties - Blocked Access

2. Select an access policy from the **Access** drop down list.
3. Click **Change**. A confirmation message dialog appears. Click **Yes** to change the access policy. The access policy gets changed from *Blocked* to the policy that you selected.

4.7 Selecting SafeStore Encryption Services Security Options

4.7.1 Enabling the Security Key Identifier, Security Key, and Password

The SafeStore Encryption Services feature provides the ability to encrypt data and use disk-based key management for the data security solution. This solution protects your data in case of theft or loss of physical drives. This section describes how to enable, change, or disable the drive security settings, and how to import a foreign configuration.

Perform the following steps to enable the encryption settings for the security key identifier, security key, and password.

1. Click **Drive Security** on the main WebBIOS dialog.

The **Drive Security** dialog appears, as shown in the following figure.



Figure 76: Driver Security Settings Dialog

2. To enable the drive security settings, select the **Enable drive security** radio button, and click **Accept**.

The Drive Security Choose mode Dialog wizard appears as shown in the following figure. This wizard allows you to choose the security key management mode.

- External Key Management (EKM) mode
- Local Key Management (LKM) mode

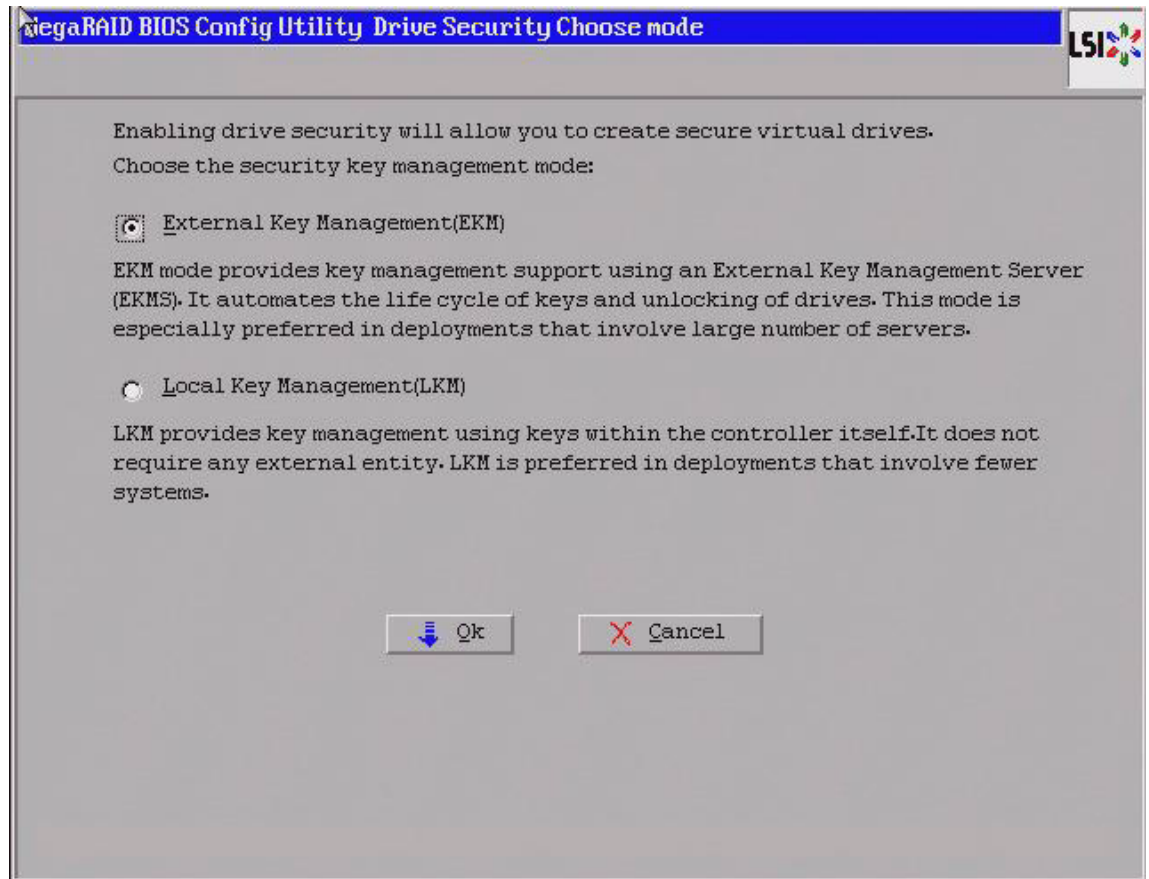


Figure 77: Drive Security Choose Mode Dialog

3. Select **EKM** in the Drive Security Choose mode wizard.
4. Click **OK**. The system asks you for a restart. The restart message appears as shown in [Figure 82](#).

If you select **LKM** in **Drive Security Choose Mode** dialog ([Figure 77](#)), a security key identifier is created, as shown in the following figure.

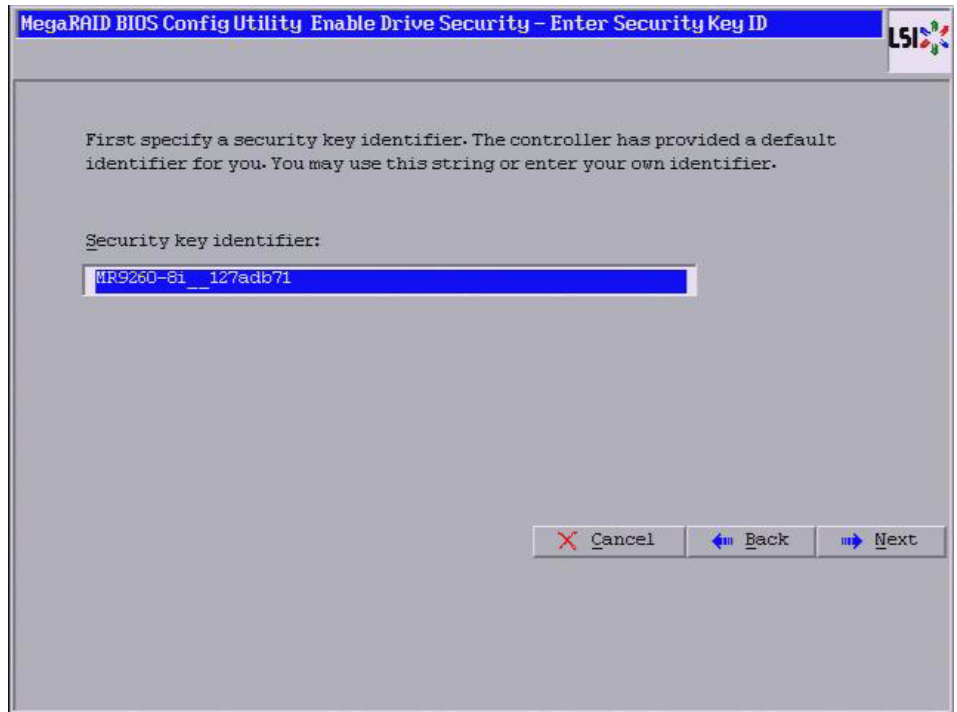


Figure 78: Enable Drive Security – Enter Security Key ID Dialog

5. Either accept the default security key ID, or enter a new security key ID.
6. Click **Next**.

The Enable Drive Security - Enter Security Key Dialog appears, as shown in [Figure 79](#).




Figure 79: Enable Drive Security – Enter Security Key

7. Either enter a new drive security key, or click **Suggest** to fill the new security key. Enter the new drive security key again to confirm.

The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.

8. Click **Next**.

The Enable Drive Security – Enter Pass Phrase dialog appears as shown in [Figure 80](#). You have the option to provide a pass phrase for additional security.

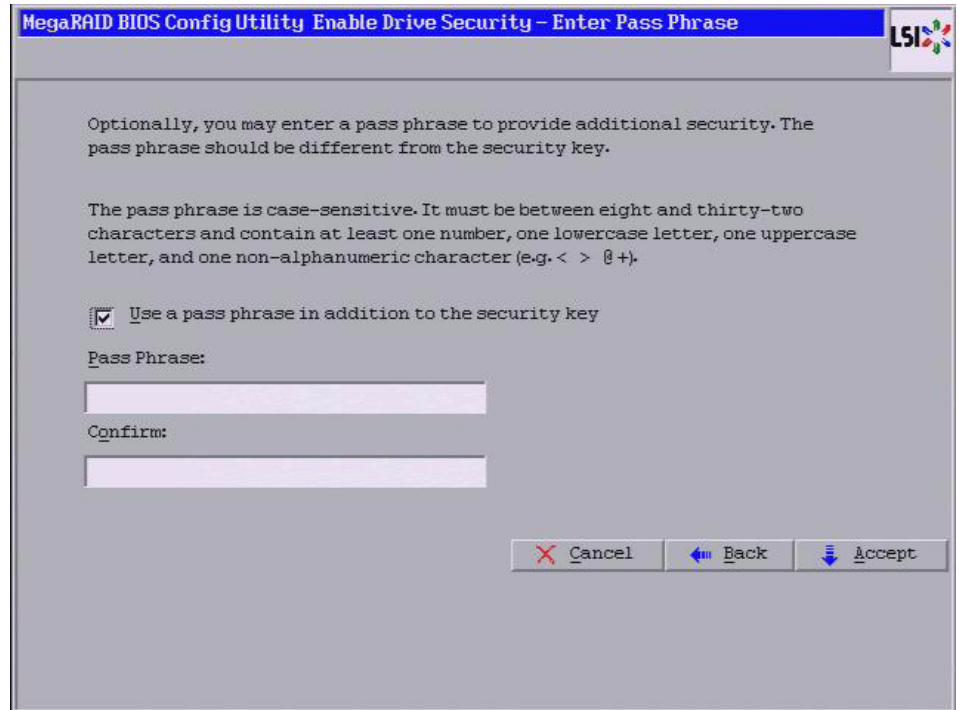


Figure 80: Enable Drive Security – Enter Pass Phrase

9. To use a pass phrase, select the **Use a pass phrase in addition to the security key** check box.
10. Enter a new pass phrase, and enter the new pass phrase again to confirm.

The pass phrase is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.

NOTE: Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the **Pass Phrase** field or **Security Key** field. Firmware works only with the ASCII character set.

11. Click **Accept**.

The **Confirm Enable Drive Security** dialog appears, as shown in [Figure 81](#).



Figure 81: Confirm Enable Drive Security Dialog

12. Click **Yes** to confirm that you want to enable the drive security settings.

WebBIOS enables the security key ID, the security key, and the pass phrase (if applicable) that you entered and returns you to the WebBIOS main menu.

CAUTION: If you forget the security key, you will lose access to your data. Be sure to record your security key information. You might need to enter the security key to perform certain operations.

4.7.2 Enabling Drive Security using EKM

EKM is used for key management when a large number of systems are deployed. You can automate and manage the life cycle of keys, and also use them to unlock configurations.

When you select EKM in the dialog shown in [Figure 77](#), and click **OK**, the following dialog appears.

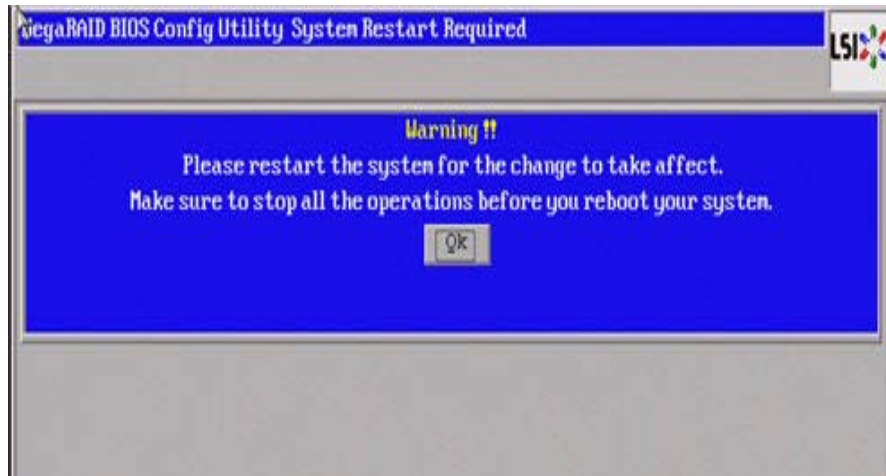


Figure 82: System Restart Request Dialog

4.7.3 Changing the Security Key Identifier, Security Key, and Pass Phrase

If you selected disk-based encryption when you made the RAID configuration, the drive security is enabled. Perform the following steps to change the encryption settings for the security key identifier, security key, and pass phrase.

1. Click **Drive Security** in the BIOS Config Utility window.

The **Change Drive Security** dialog appears as shown in [Figure 83](#).



Figure 83: Drive Security Settings Dialog

2. To change the drive security settings, select the **Change drive security settings** radio button and click **Accept**.

The Change Security Settings – Introduction dialog appears as shown in [Figure 84](#).



Figure 84: Config Utility Change Security Settings Dialog

4.7.4 Change Security from EKM to LKM

1. Select either one of the radio buttons:
 - Select the **Change current security settings** radio button, and click **OK**.
The system restart message appears, as shown in [Figure 82](#).
 - Select the **Switch to Local Key management (LKM) mode** radio button.
The Change Security Settings – Security Key ID dialog appears, as shown in [Figure 85](#).

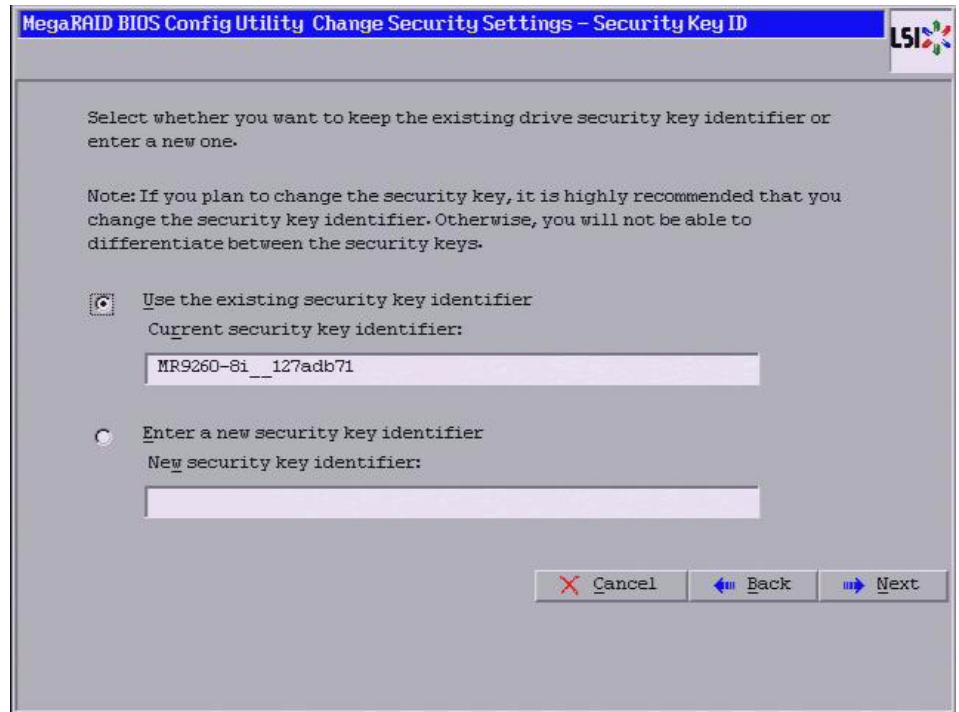


Figure 85: Change Security Settings – Security Key ID Dialog

2. Choose whether to use the existing security key ID or enter a new security key ID. You have the following options:
 - Use the existing security key identifier (Current security key identifier).
 - Enter a new security key identifier (New security key identifier).
3. Click **Next**.

The Change Security Settings – Security Key dialog appears, as shown in [Figure 86](#). You have the option to either use the existing security key or enter a new one.

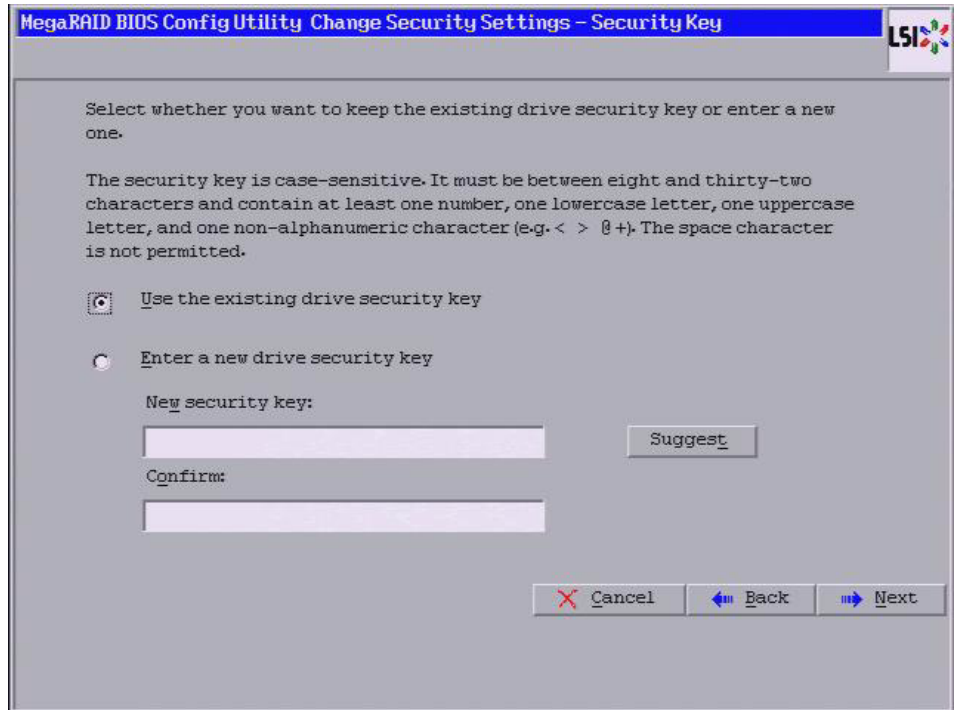


Figure 86: Change Security Settings – Security Key Dialog

4. To create a new drive security key, either enter a new drive security key in the **New security key** field, or click **Suggest** to fill the new security key.

The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.

5. Enter the new drive security key again in the **Confirm** field.
6. Click **Next**.

The Change Security Settings – Pass Phrase dialog appears, as shown in [Figure 87](#).

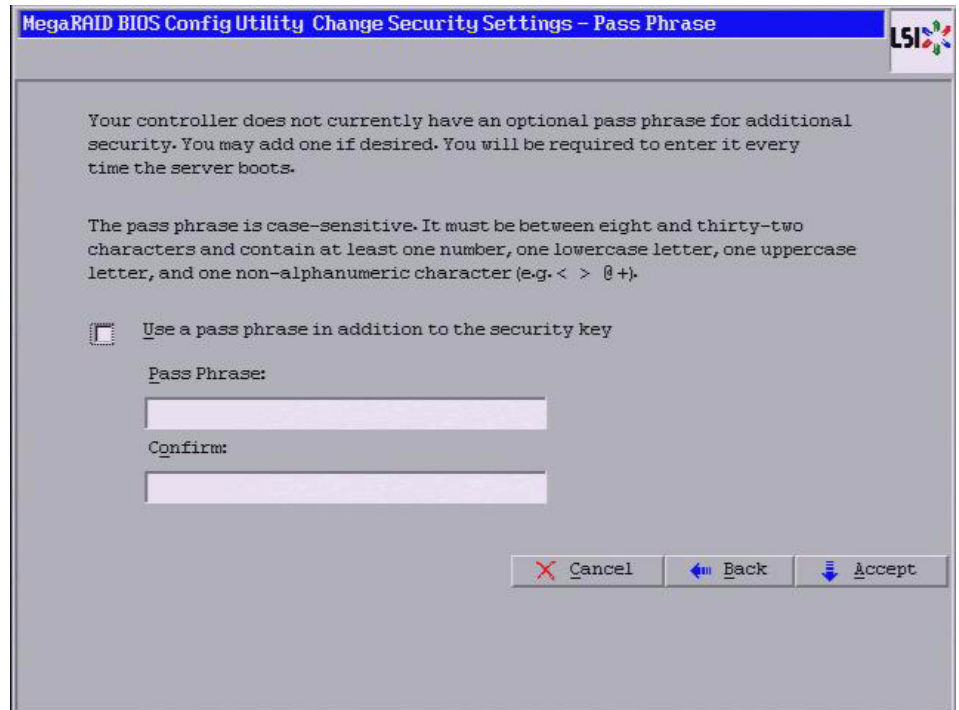


Figure 87: Change Security Settings – Pass Phrase Dialog

7. To use a pass phrase, click the **Use a pass phrase in addition to the security key** check box.

8. Enter a new pass phrase, and enter the new pass phrase again to confirm.

The pass phrase is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.

Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the **Pass Phrase** field or **Security Key** field. The firmware works only with the ASCII character set.

9. Click **Accept**. The Authenticate Drive Security Settings dialog (Figure 88) appears in either of the two scenarios mentioned below.
 - If you entered a new pass phrase.
 - If you entered a new drive security key.

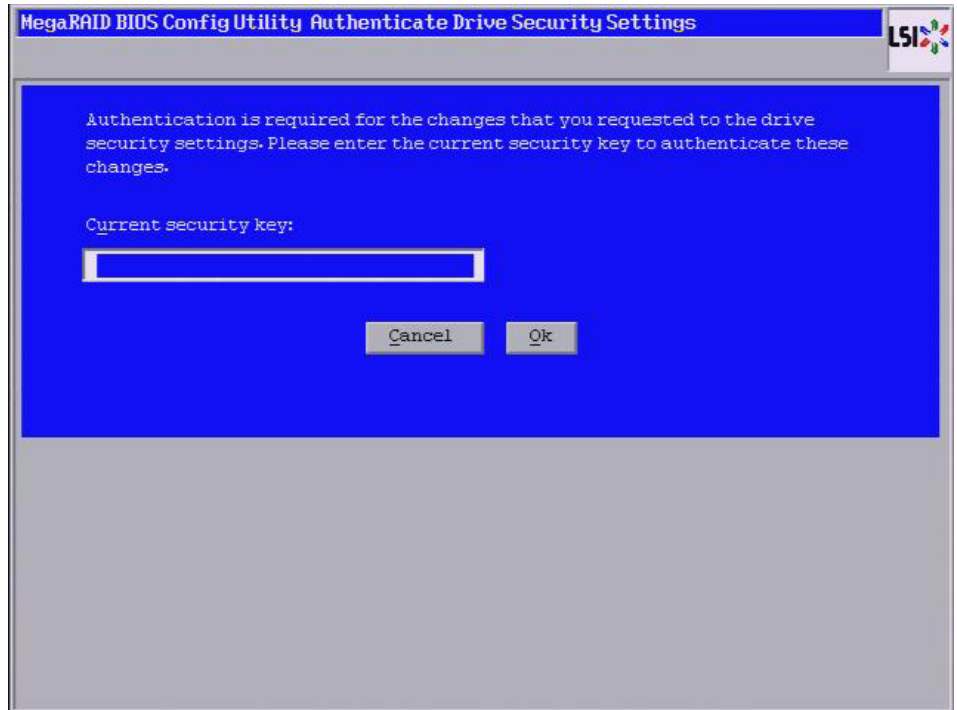


Figure 88: Authenticate Drive Security Settings Dialog

10. Enter the current security key, and click **OK**.

The text box for the security key can hold up to 32 characters. The key must be at least eight characters. After you enter the correct security key, the Confirm dialog appears, as shown in [Figure 89](#).




Figure 89: Confirm Page Dialog

11. Click **Yes** to confirm that you want to change the drive security settings.

If the current security key is not needed, WebBIOS saves the changes to the security settings and returns you to the main menu. If the current security key is needed, the Authenticate Drive Security Settings dialog appears.

4.7.5 Changing Security from LKM to EKM

Perform the following steps if you want to change security from LKM to EKM.

1. Select **Drive Security** from the WebBIOS main screen.
2. Select **Change drive security**.

The **Change Security Settings** wizard appears, as shown in the following figure.

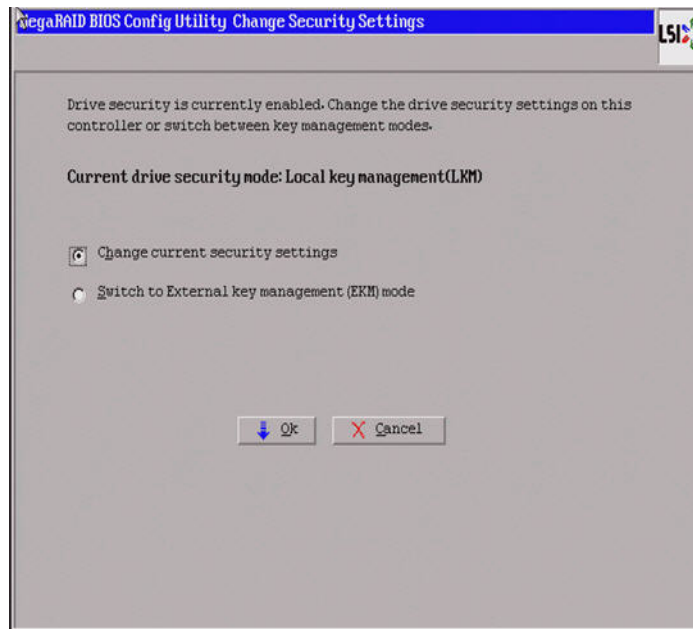


Figure 90: Change Security Settings

3. Select the **Change current security settings** radio button, and click **Ok**.

The **Change Security Settings – Security Key ID** dialog is displayed (Figure 85). This procedure continues until the step that contains Figure 89.

4. Select the **Switch to External Key Management (EKM) mode** radio button and click **Ok**.

The Authenticate Drive Security Settings dialog is displayed (Figure 88).

4.7.6 Disabling the Drive Security Settings

Perform the following steps to disable the drive security settings.

NOTE: If you disable the drive security settings, you cannot create any new secure virtual drives. Disabling these settings does not affect the security or data of foreign drives. If you removed any drives that were previously secured, you must enter the security key when you import settings.

1. Click **Drive Security** on the main WebBIOS dialog.

The Drive Security dialog appears, as shown in Figure 91.



Figure 91: Drive Security Dialog

2. To disable the drive security settings, select the **Disable drive security** radio button and click **Accept**.

The Confirm Disable Drive Security page dialog appears, as shown in [Figure 92](#).

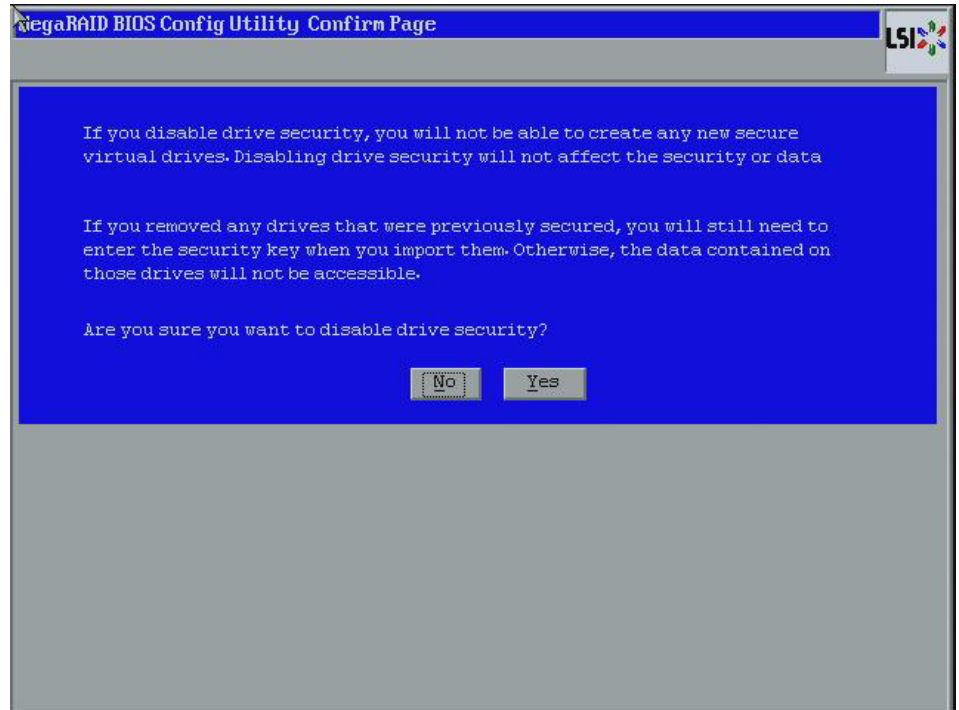


Figure 92: Confirm Disable Drive Security Settings Dialog

3. Click **Yes** to confirm that you want to disable the drive security settings.

WebBIOS returns you to the main menu.

4.8 Viewing and Changing Device Properties

4.8.1 Viewing Controller Properties

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

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 WebBIOS dialog. When the Adapter Selection dialog appears, select the controller you want from the list.

Follow these steps to view the properties of the currently selected controller.

1. Click **Controller Properties** on the main WebBIOS dialog.

There are four Controller Information dialogs. [Figure 93](#) shows the first dialog.



Figure 93: First Controller Information Dialog

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

NOTE: If you are using CacheCade Pro 2.0, four additional fields will appear in the Controller Properties screen - **CacheCade SSD Caching**, **Write Cache Capable**, **Total Cache Size**, and **Maximum Cache Size**.

2. Click **Next** to view the second Controller information dialog, as shown in Figure 94.



Figure 94: Second Controller Information Dialog

3. Click **Next** to view the third Controller information dialog, as shown in Figure 95.

Properties			
Battery Backup	Present	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Disabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	4
Reconstruction Rate	30	Spinup Delay	2
NCQ	Enabled		

Submit Reset Next

Home Back

Figure 95: Third Controller Properties Dialog

NOTE: If you are using CacheCade Pro 2.0, an additional field, **SSD Caching** will appear in the Controller Properties screen.

4. Click **Next** to view the fourth Controller Information dialog, as shown in [Figure 96](#).

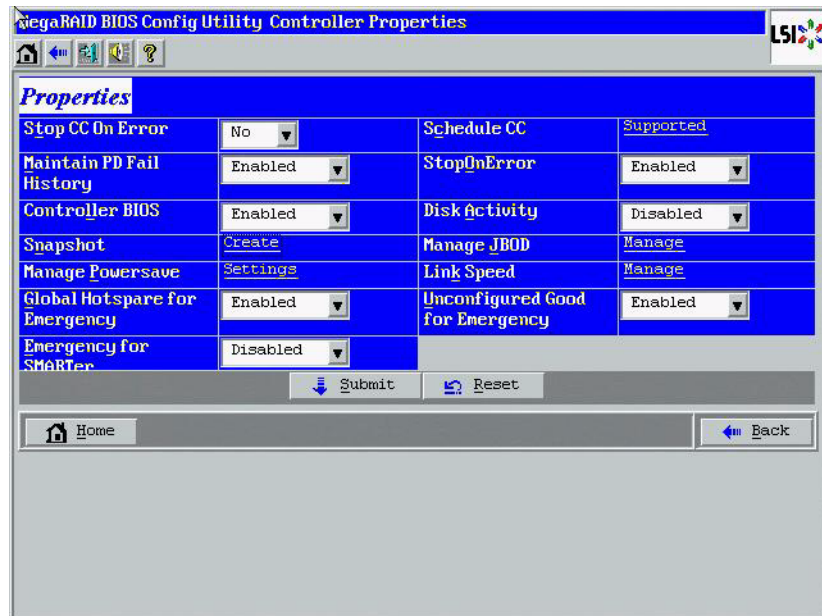


Figure 96: Fourth Controller Properties Dialog

[Table 21](#) describes the entries and options listed on the second and third Controller Information dialog. 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 21: Controller Information Menu Options

Option	Description
Battery Backup	This entry indicates whether the selected controller has a BBU. If present, you can click Present to view information about the BBU. For more information, see Section 4.8.5, Viewing and Changing Battery Backup Unit Information
Set Factory Defaults	Use this option to load the default MegaRAID WebBIOS configuration utility settings. The default is No .
Cluster Mode	Use this option to enable or disable Cluster mode. The default is Disabled . 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 that are devoted to a rebuild.
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.

Table 21: Controller Information Menu Options (Continued)

Option	Description
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 Enabled. 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 Enabled.
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 None, 128MB-way, and 1GB-way. The default is 1GB-way. 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 (self-monitoring analysis and reporting technology [SMART] 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 Enabled.
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.
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 4 drives.
Spinup Delay	Use this option to control the interval (in seconds) between spin up 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.
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 Enabled.
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 No.
Maintain PD Fail History	Enable this option to maintain the history of all drive failures. The default is Enabled.
Schedule CC	Indicates whether the option to schedule the date and time for a consistency check is supported.

Table 21: Controller Information Menu Options (Continued)

Option	Description
Snapshot	Use this option to create a snapshot of a volume. MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if data is deleted accidentally or maliciously, restore the data from the view or roll back to a snapshot at a previous point-in-time (PIT). MegaRAID Recovery supports up to eight snapshots of PiTs for each volume.
Disk Activity	Enable this property if you want to locate a particular disk. This disk can be identified with a continuous blinking of green activity LED. This works only if the disks are installed in an enclosure.
Manage JBOD	Converting the multiple JBOD drives to unconfigured drive at single selection.
Global hot spare for Emergency	Use this option to specify if it is acceptable to commission incompatible global hot spare drives as emergency hot spare drives.
Unconfigured good for Emergency	Use this option to specify if the commission unconfigured good drives can be used as emergency hot spare drives.
Emergency for SMARTer	Use this option to specify if it is acceptable to commission emergency hot spare drives for predictive failure analysis(PFA) events.
Manage Powersave	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured drives, hot spares, and configured drives.
Link Speed	Use this option to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

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

4.8.2 Viewing Virtual Drive Properties, Policies, and Operations

WebBIOS displays properties, policies, and operations for virtual drives.

To view these items for the currently selected virtual drive, click on a virtual drive icon in the right panel on the WebBIOS Configuration utility main dialog.

The Virtual Drive dialog appears, as shown in [Figure 97](#).

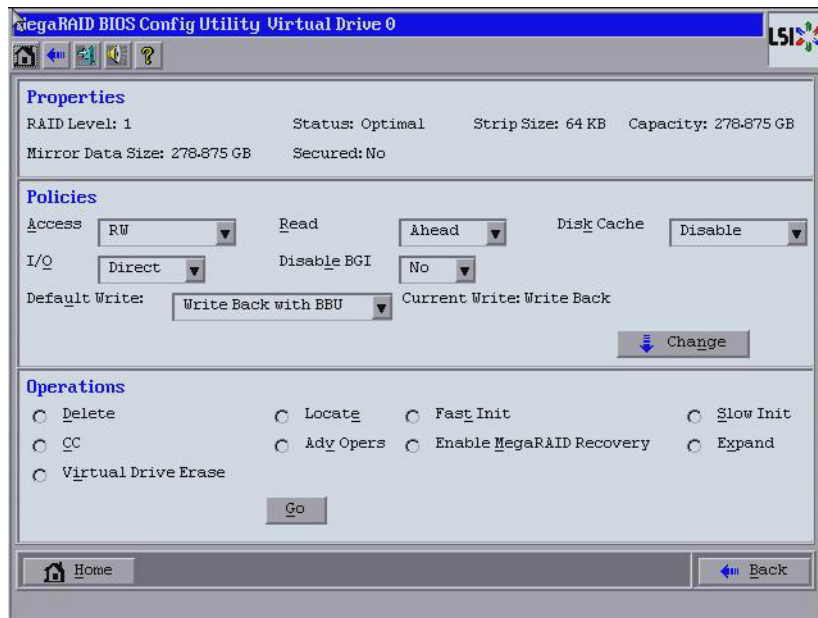


Figure 97: Virtual Drive Dialog

The Properties panel of this dialog displays the virtual drive's RAID level, state, capacity, strip size, and metadata 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 4.5.2, Using Manual Configuration](#). To change any of these policies, make a selection from the drop-down list 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**. Choose from the following options:

- Select **Delete** to delete this virtual drive. For more information, see [Section 4.14.2, Deleting a Virtual Drive](#).
- Select **Locate** to make the LEDs blink on the drives used by this virtual drive. This action works only if the drives are installed in a drive enclosure that supports SCSI-Accessed-Fault-Tolerant-Enclosure (SAFTE).
- Select **Fast Init** or **Slow Init** to initialize this virtual drive. A fast initialization quickly writes zeroes to the first and last 10-MB 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 the drive.

- Select **CC** to run a consistency check on this virtual drive. For more information, see [Section 4.14.1, *Running a Consistency Check*](#). (This option is not available for RAID 0 virtual drives.)

- Select **Adv Opers** to access dialogs to remove drives, migrate RAID levels (that is, change the virtual drive configuration by adding a drive and changing the RAID level), and use MegaRAID Recovery.

See [Section 4.14.8, *Migrating the RAID Level of a Virtual Drive*](#) for information about adding a drive to a virtual drive or migrating its RAID level. See [Section 4.11, *Using MegaRAID Recovery*](#) for the MegaRAID Recovery procedure.

- Select **Enable MegaRAID Recovery** to use MegaRAID Recovery, also known as Snapshot. Recovery offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file.

See [Section 4.11, *Using MegaRAID Recovery*](#) for the MegaRAID Recovery procedure.

- Select **Expand** to increase the size of a virtual drive to occupy the remaining capacity in the drive group.

See [Section 4.9, *Viewing and Expanding a Virtual Drive*](#) for the procedure you can use to expand a virtual drive.

- Select **Virtual Drive Erase** to erase a virtual drive and over write all users.

See [Section 4.12.2, *Virtual Drive Erase*](#) for the procedure you can use to erase a virtual drive.

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

4.8.3 Viewing Drive Properties

The Physical Drive dialog 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 dialog:

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

Figure 98 shows the Physical Drive dialog.

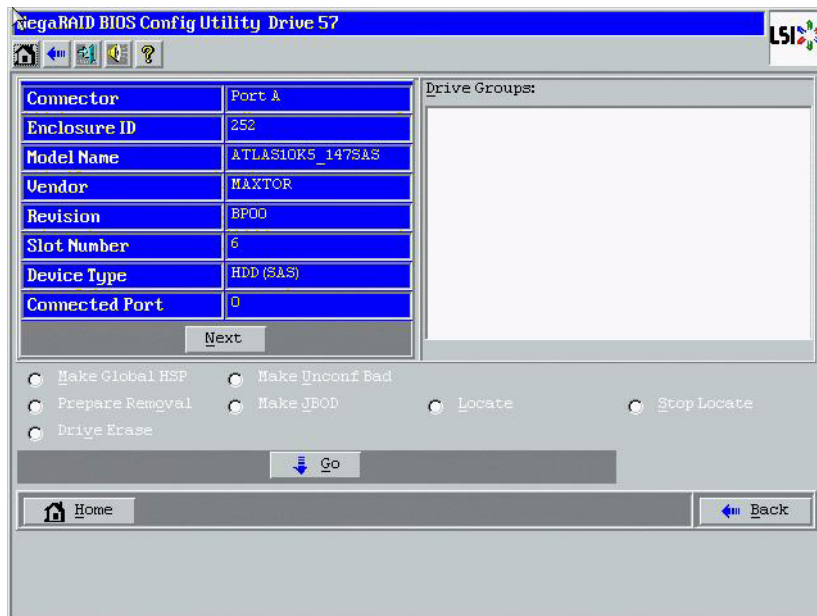


Figure 98: Physical Drive Dialog

The drive properties are read-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 dialog. 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 operation works only if the drive is installed in a drive enclosure.

If the drive state is Unconfigured Good, the following additional operations appear on this dialog.

- Select **Make Global HSP** to make a global hot spare, which is 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 appears under its associated virtual drive. The drive number, drive state, drive capacity, and drive manufacturer appear.

- Select **Enclosure Affinity** so drive failures are present on a split backplane configuration, then the hot spare will be used first on the backplane side in which it resides.
- 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 power save mode because it also involves flagging the drive as ready to remove. Therefore, if you choose to prepare a drive for removal, selecting **Ready to Remove** displays in the device tree for that drive, instead of **Powersave**.
- Select **Stop Locate** to stop the LED flash on the drive. This works only if the drive is installed in a drive enclosure.
- Select **Drive Erase** to securely erase data on non self-encrypting drives (Non-SED), which are normal HDDs.

4.8.4 Shield State

Physical devices in MegaRAID firmware transit between different states. If the firmware detects a problem or a communication loss for a physical drive, the firmware transitions the drive to a bad (FAILED or UNCONF BAD) state. To avoid transient failures, an interim state called the shield state is introduced before marking the drive as being in a bad state.

The shield state is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostic tests fail, the physical drive transitions to a bad state (FAILED or UNCONF BAD).

4.8.4.1 Shield State Physical View

Follow these steps to check if a physical drive is in a Shield state in the Physical view.

1. Click **Physical View** in the main dialog.

The physical drive that is in a shield state is marked as Shielded.

The physical view shield state is shown in [Figure 99](#).

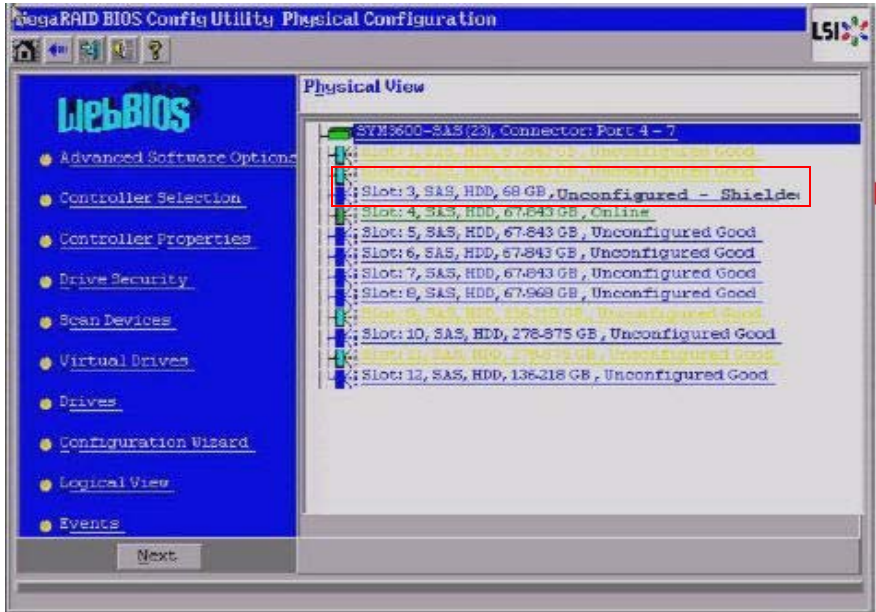


Figure 99: Physical View Shield State Dialog

4.8.4.2 Logical View Shield State

Follow these steps to view the Shield state in the Logical view.

1. Click **Logical View** in the main page.

The physical drive that is in a shield state is marked as Shielded.

The Logical view shield state is shown in [Figure 100](#).

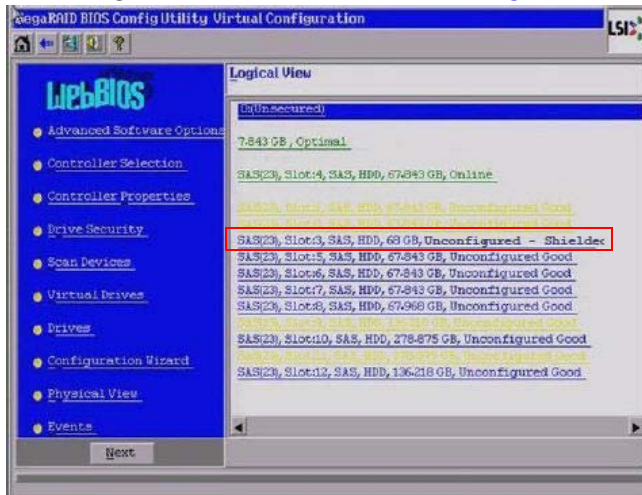


Figure 100: Logical View Shield State

4.8.4.3 Viewing the Physical Drive Properties of a Drive In Shield State

Follow these steps to view the physical properties of the drive in Shield state.

1. Click on the **Physical view** tab or the **Logical view** tab in the device tree.
2. Click the physical drive that is in shield state on the physical or logical view of device tree to view the properties.

The device properties of the drive are displayed as shown in [Figure 101](#).

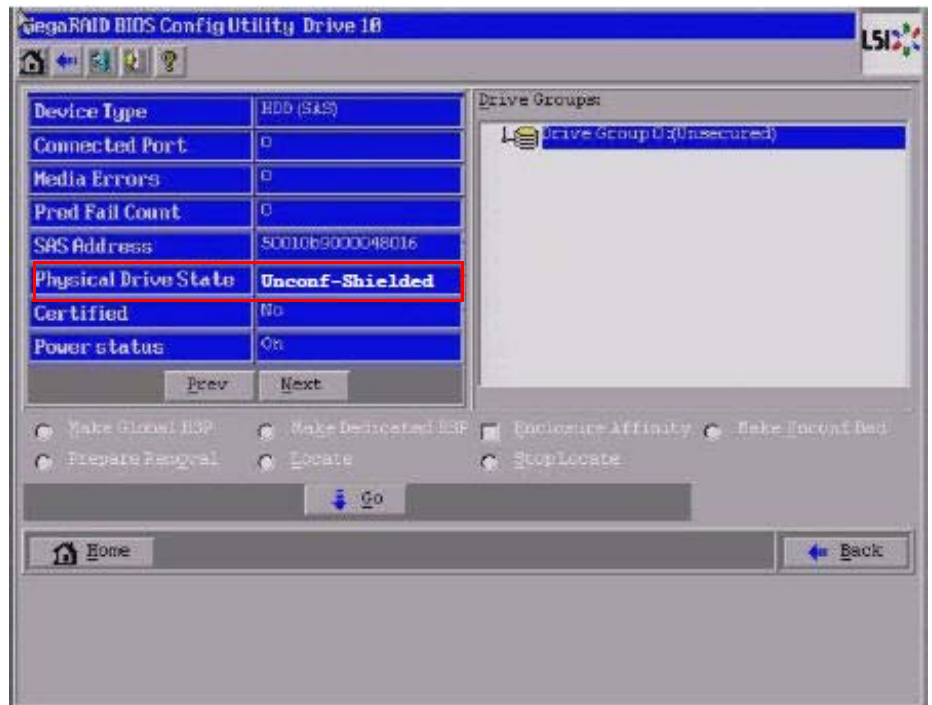


Figure 101: Physical Drive Properties of a Drive in Shield State

4.8.4.4 Viewing if Shield State is Enabled in a Controller

Follow these steps to check if the Shield state is enabled in a controller.

1. Click **Controller Properties** on the WebBIOS main menu.

- The Shield State Supported Column is displayed, as shown in Figure 102.




Figure 102: Shield State Support

4.8.5 Viewing and Changing Battery Backup Unit Information

If your SAS controller has a battery backup unit (BBU), you can view information about it and change some settings. To perform these tasks, follow these steps:

- Click **Controller Properties** on the WebBIOS main screen.

The first Config Utility Controller Information dialog appears, as shown in Figure 103.

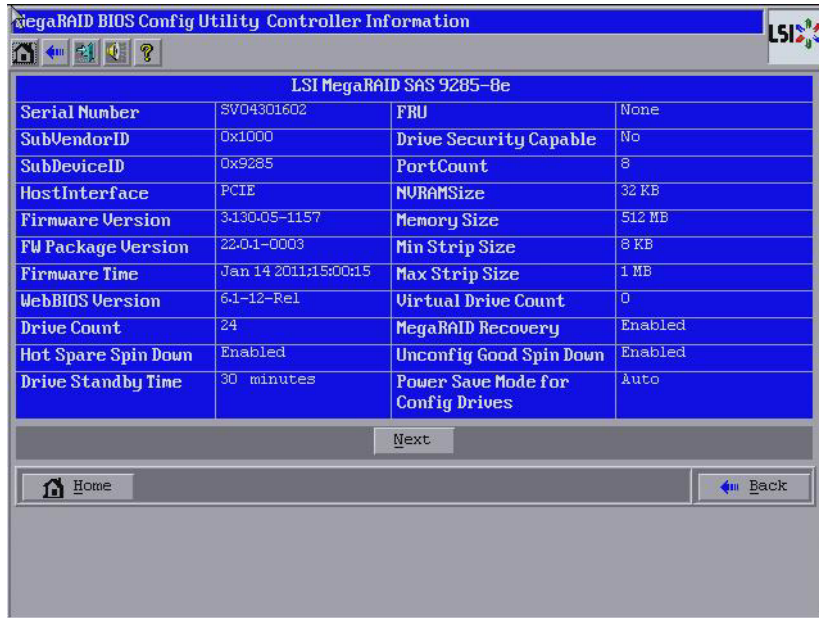


Figure 103: First Config Utility Controller Information

2. Click **Next**. The second Config Utility Controller Information dialog appears.
3. Click **Next** to view the third Config Utility Controller Properties dialog.

The third Config Utility Controller Properties dialog appears, as shown in [Figure 104](#). The **Battery Backup** field at the top-left of the dialog indicates whether the iBBU is present.

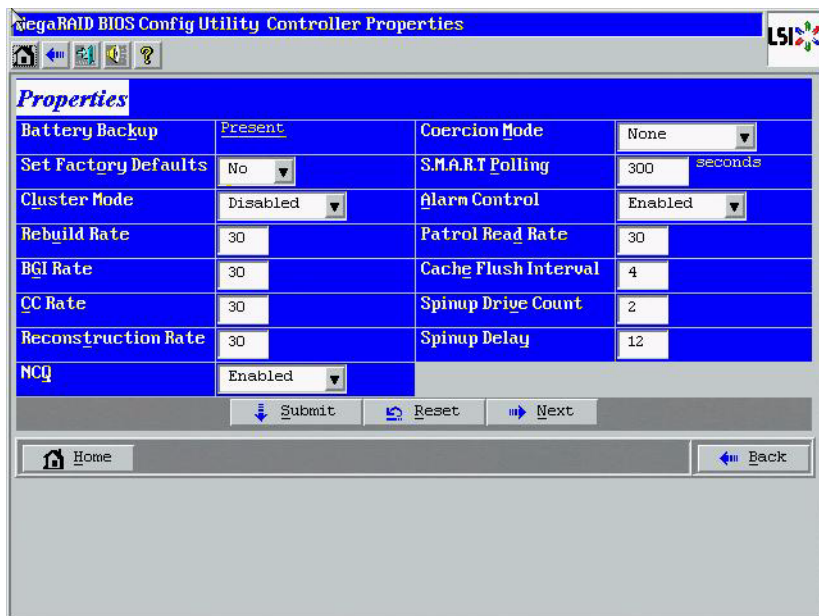


Figure 104: Third Config Utility Controller Properties

4. Click **Present** in the **Battery Backup** field.

The Config Utility Battery Module dialog appears, as shown in [Figure 105](#) and [Figure 106](#). This dialog contains the following information:

- Battery information
- Battery state
- Battery voltage
- Battery current
- Battery temperature
- Battery replacement required
- BBU retention time
- Time to full charge
- Design information
- Capacity information
- Auto learn properties and settings

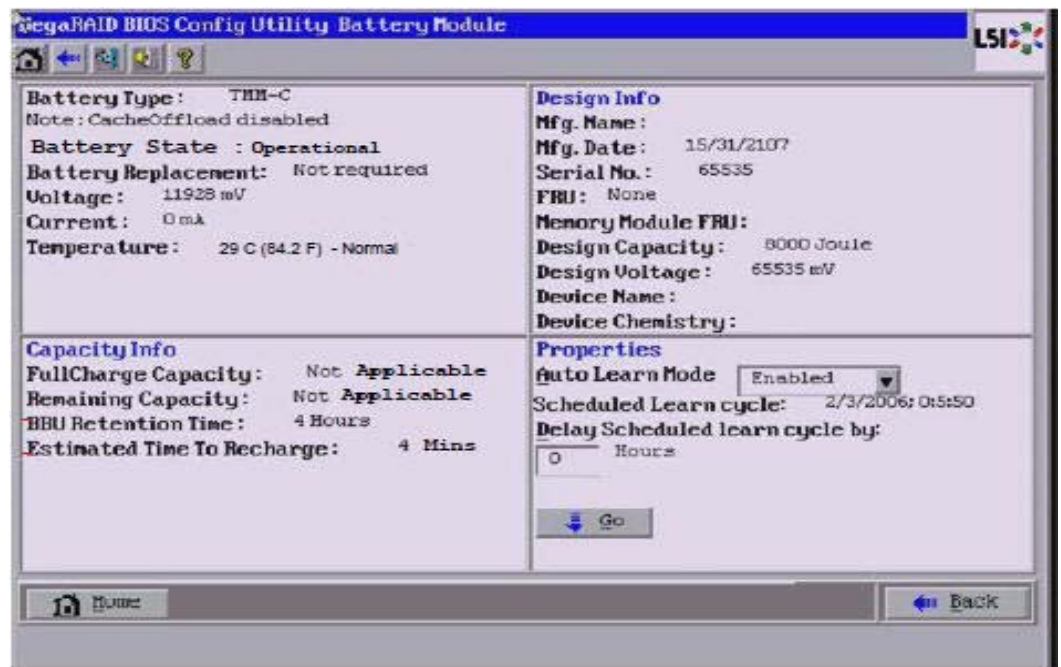


Figure 105: Config Utility Battery Module Dialog for a TMM-C Battery

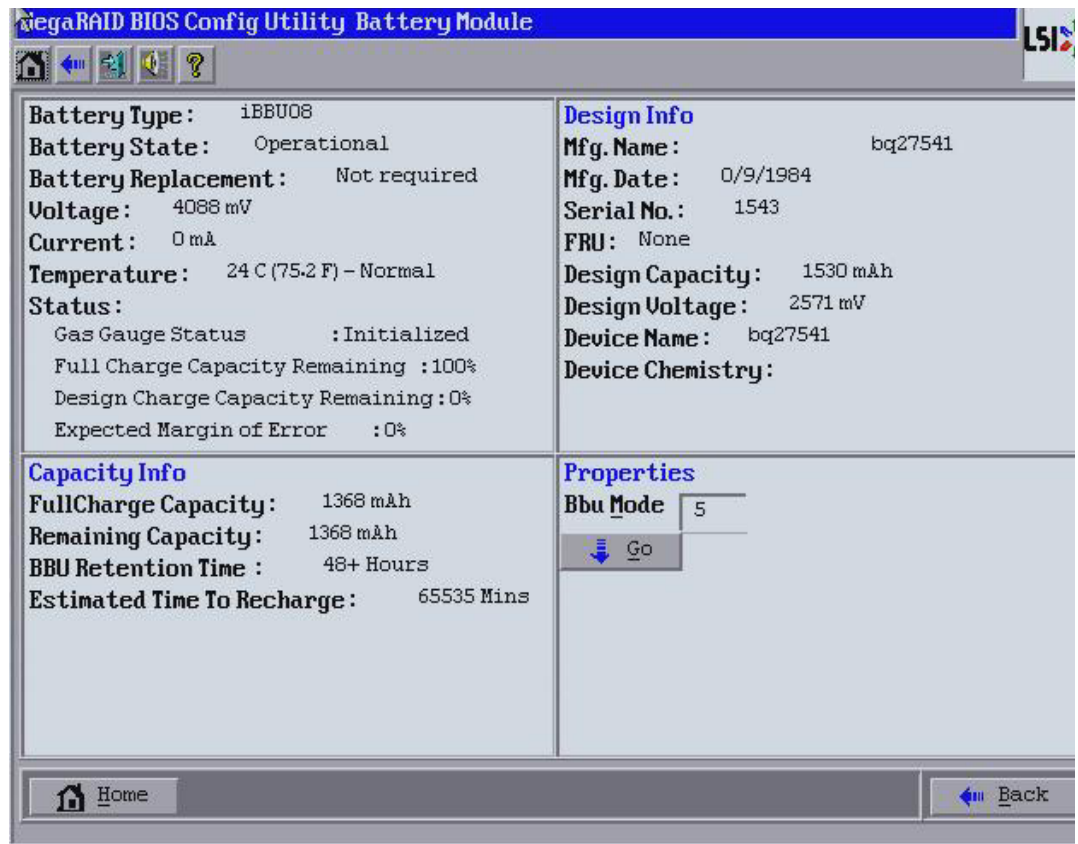


Figure 106: Config Utility Battery Module Dialog for an iBBU Battery

Most of the battery module properties are view-only and are self-explanatory. The following table gives a description of each of the BBU modes.

Table 22: BBU Modes

Mode of Operation MR_MFC_DEFAULTS.bbu Mode	Description
0	48 hours retention @ 60oC, 1 year Service Life
1	12 hours retention @ 45oC, 5 year Service Life, transparent learn
2	12 hours retention @ 55oC, 3 year Service Life, transparent learn
3	24 hours retention @ 45oC, 3 year Service Life, transparent learn
4	48 hours retention @ 45oC, 3 year Service Life
5	48 hours retention @ 55oC, 1 year Service Life
6	Same as the description for BBU mode 5. The BBU mode 6 enables you to get events when the battery capacity reaches sub-optimal and critical thresholds.

The lower right corner of the dialog (Figure 105) contains the auto learn options. A *learning cycle* is a battery calibration operation performed by the controller periodically to determine the condition of the battery. You can change the learn delay interval (the length of time between automatic learning cycles) and the auto learn mode.

NOTE: LSI recommends leaving the learn delay interval and the auto learn mode at their default settings.

NOTE: After selecting **Disabled** (in the **Auto Learn Mode**), if you select **Enabled**, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. However, in the **Scheduled Learn cycle** field, the value **None** is displayed. The **Scheduled Learn cycle** field is not updated until the battery relearn is completed. Once the relearning cycle is completed, the value in the **Scheduled Learn cycle** field displays the new date and time of the next battery learning cycle.

4.8.5.1 Setting the Learn Delay Interval

The learn delay interval is the length of time between automatic learning cycles. Perform the following steps to change the interval:

- a. Open the drop-down list in the **Auto Learn Mode** field.
- b. Select the learn mode as **Enabled** (the default).
Use this setting so that the controller performs the learning cycle automatically.
- c. Change the number of hours in the **Delay Scheduled learn cycle by** field.
You can delay the start of the learn cycles for up to 168 hours (7 days).
- d. Click **Go** to set the interval.

4.8.5.2 Setting the Auto Learn Mode

You can start battery learning cycles manually or automatically. The Auto Learn modes are:

- **BBU Auto Learn:** The firmware tracks the time since the last learning cycle and performs a learn cycle when due.
- **BBU Auto Learn Disabled:** The firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.

NOTE: After selecting **Disabled** (in the **Auto Learn Mode**), if you select **Enabled**, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. However, in the **Scheduled Learn cycle** field, the value **None** is displayed. The **Scheduled Learn cycle** field will not be updated until the battery relearn is completed. Once the relearning cycle is completed, the value in the **Scheduled Learn cycle** field will display the new date and time of the next battery learning cycle.

- **BBU Auto Learn Warn:** The firmware warns about a pending learning cycle. You can initiate a learning cycle manually. After the learning cycle is complete, the firmware resets the counter and warns you when the next learning cycle time is reached.

Perform the following steps to choose an auto learn mode:

1. Open the drop-down list in the **Auto Learn Mode** field.
2. Select an auto learn mode.
3. Click **Go** to set the auto learn mode.

NOTE: When you replace the iBBU, the charge cycle counter is reset automatically.

4.8.6 Managing Link Speed

The Managing Link Speed feature allows you to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

All phys in a SAS port can have different link speeds or can have the same link speed.

You can select a link speed setting. However, if phys in a SAS port have different link speed settings and if a phy is connected to a drive or an expander, the firmware overrides the link speed setting you have selected and instead uses the common maximum link speed among all the phys.

To change the link speed, perform the following steps:

1. Click **Controller Properties** on the WebBIOS main menu.

There are four Controller Properties screens. The first Controller Properties screen appears. See [Figure 93](#) to view this screen.

2. Click **Next** to access the second Controller Properties screen.

The second Controller Properties screen appears. See [Figure 94](#) to view this screen.

3. Click **Next** to access the third Controller Properties screen.

The third Controller Properties screen appears. See [Figure 95](#) to view this screen.

4. Click **Next** to access the fourth Controller Properties screen.

The fourth Controller Properties screen appears, as shown in the following figure.

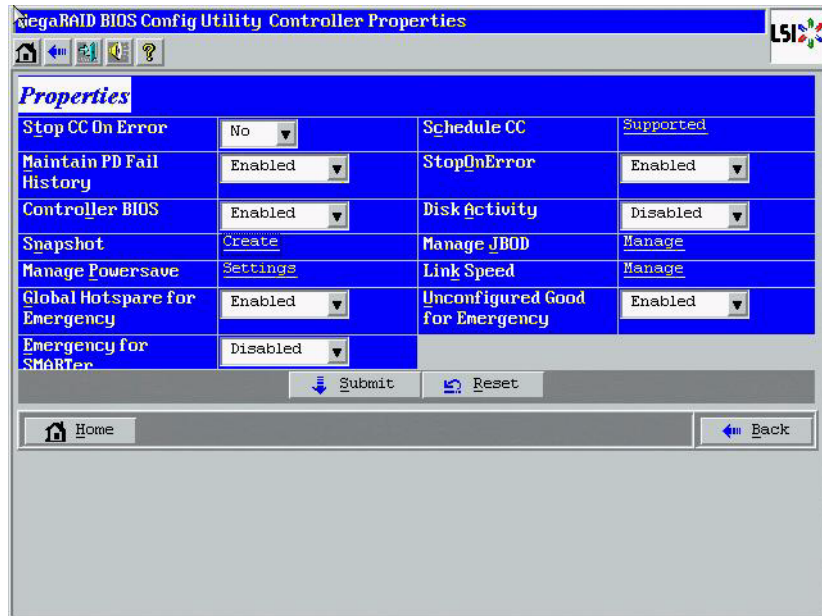


Figure 107: Fourth Controller Properties Screen

5. Click **Manage** in the **Link Speed** field.

The Manage Link Speed dialog box appears, as shown in the following figure.

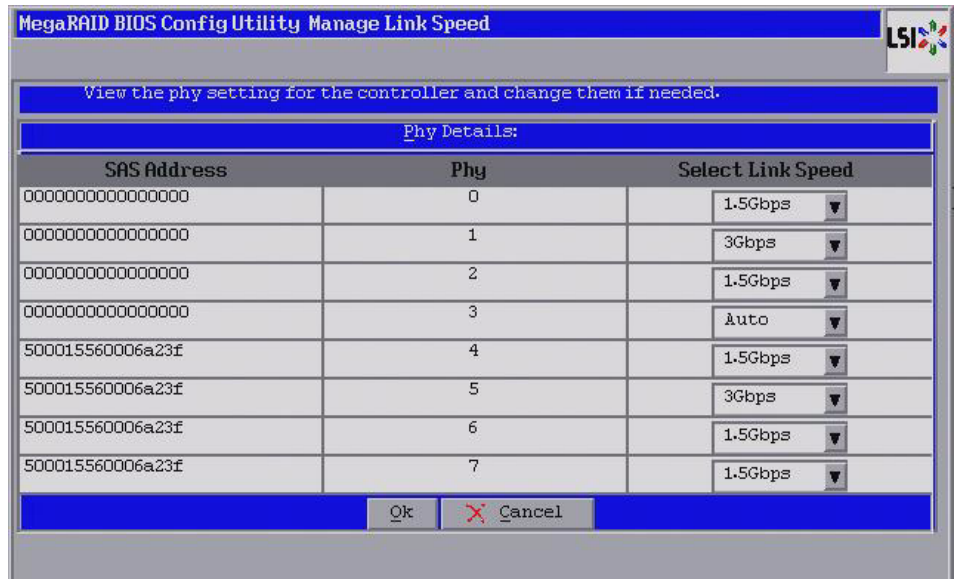


Figure 108: Manage Link Speed Screen

- The SAS Address column displays the SAS address that uniquely identifies a device in the SAS domain.
- The Phy column displays the system-supported phy link values. The phy link values are from 0 through 7.
- The Select Link Speed column displays the phy link speeds.

6. Select the desired link speed from the **Select Link Speed** field using the drop-down selector.

The link speed values are Auto, 1.5, 3.0 or 6.0 Gbps.

NOTE: By default, the link speed in the controller is *Auto* or the value last saved by the user.

7. Click **OK**.

The link speed value is now reset. The change takes place after you restart the system. The message box appears, as shown in the following figure.



Figure 109: System Restart Required Message Box

4.8.7 Viewing Enclosure Properties

Using WebBIOS, you can view the enclosure properties of all of the enclosures connected to the server.

Follow these steps to view enclosure properties.

1. Go to the Physical view of the WebBIOS Utility.
2. Click the enclosure node.

The enclosure properties are displayed, as shown in Figure 110.

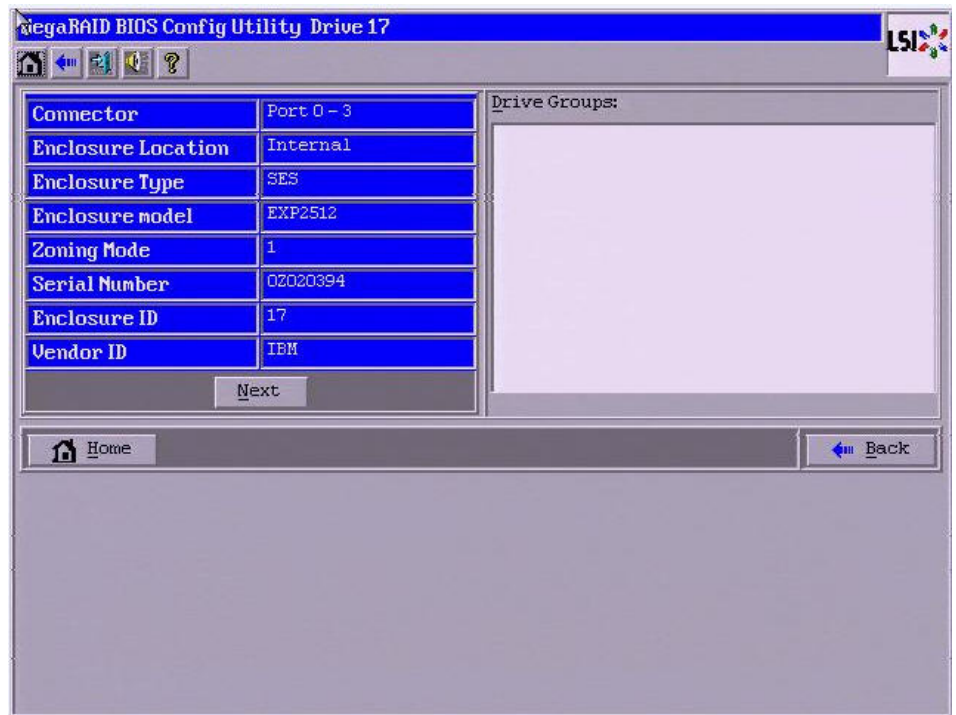


Figure 110: Enclosure Properties

3. Click **Next** to view additional properties, as shown in [Figure 111](#).

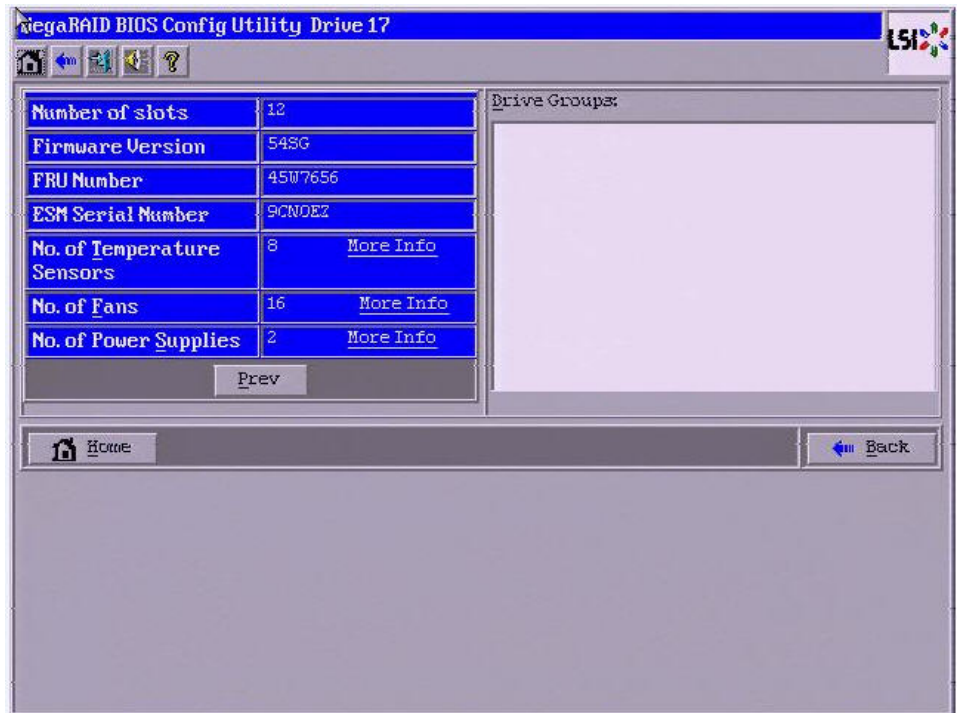
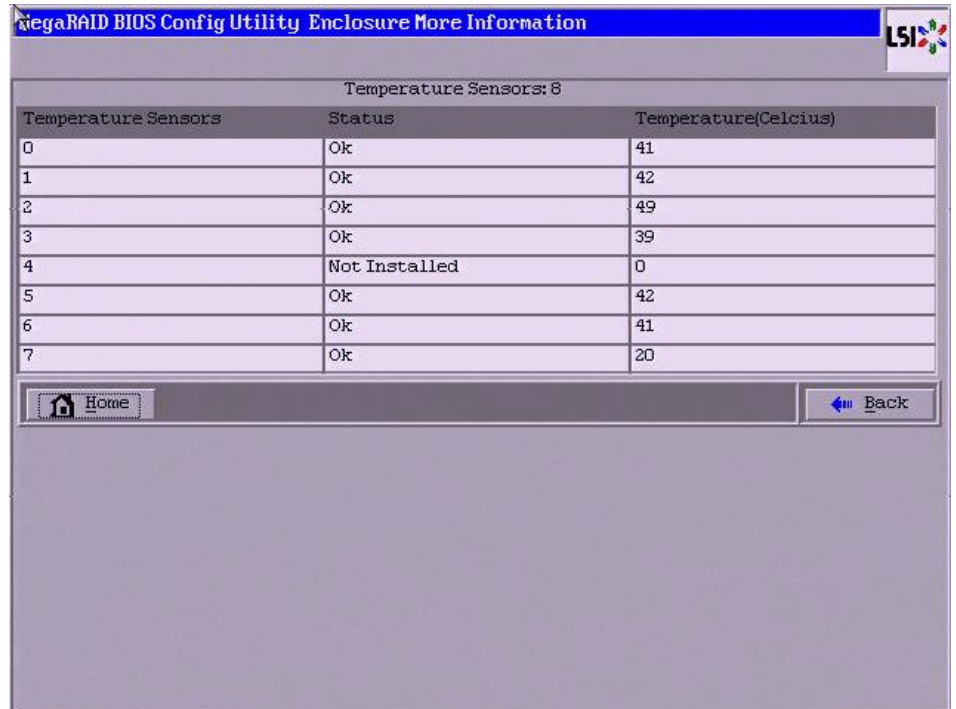


Figure 111: Additional Enclosure Properties

4. Click **More Info** to view additional information on the number of temperature sensors (Figure 112), number of fans (Figure 113), and the number of power supplies (Figure 114).



The screenshot displays the 'MegaRAID BIOS Config Utility Enclosure More Information' window. The title bar includes the LSI logo. The main content area shows 'Temperature Sensors: 8' and a table with three columns: 'Temperature Sensors', 'Status', and 'Temperature(Celcius)'. The table lists 8 sensors with their respective IDs, statuses, and temperatures. At the bottom of the window, there are 'Home' and 'Back' navigation buttons.

Temperature Sensors	Status	Temperature(Celcius)
0	Ok	41
1	Ok	42
2	Ok	49
3	Ok	39
4	Not Installed	0
5	Ok	42
6	Ok	41
7	Ok	20

Figure 112: Enclosure More Information - Temperature Sensors

MegaRAID BIOS Config Utility Enclosure More Information

Fans: 16

Fan	Status	Speed(RPM)	Speed Code
0	Ok	6500	Lowest Speed
1	Ok	9300	Lowest Speed
2	Ok	6300	Lowest Speed
3	Ok	9300	Lowest Speed
4	Ok	6300	Lowest Speed
5	Ok	9400	Lowest Speed
6	Ok	6300	Lowest Speed
7	Ok	9700	Lowest Speed
8	Ok	6500	Lowest Speed
9	Ok	9000	Lowest Speed
10	Ok	6400	Lowest Speed
11	Ok	9500	Lowest Speed
12	Ok	6400	Lowest Speed
13	Ok	9200	Lowest Speed
14	Ok	6200	Lowest Speed
15	Ok	9400	Lowest Speed

Home Back

Figure 113: Enclosure More Information - Number of Fans

MegaRAID BIOS Config Utility Enclosure More Information

Power Supply: 2

Power Supply	Status
0	Ok
1	Critical

Home Back

Figure 114: Enclosure More Information - Number of Power Supplies

4.8.8 SSD Disk Cache Policy

MegaRAID supports changes to the write-cache policy for SSD media of individual physical drives.

When SSDs are configured in a mixed disk group with HDDs, the **Physical Device Write-Cache Policy** setting of all of the participating drives is changed to match the SSD cache policy setting.

4.8.8.1 Viewing Cache Properties

Follow these steps to view the SSD Disk Cache Setting property.

1. Click the controller properties link in the main menu.
2. Click **Next** to view the controller properties with **SSD Disk Cache Setting** displayed, as shown in [Figure 115](#).

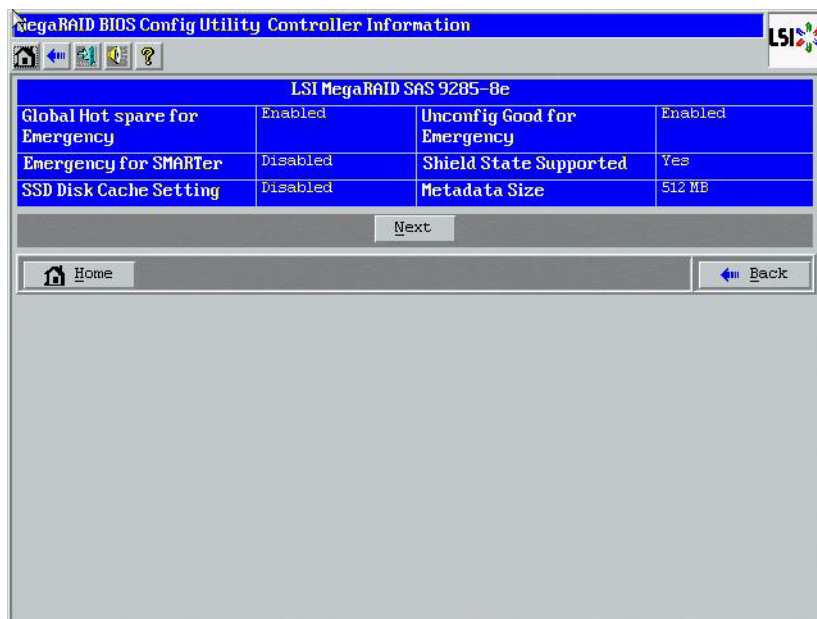


Figure 115: SSD Disk Cache Setting in Controller Properties Dialog

4.8.9 Emergency Hotspare

When a drive within a redundant virtual drive fails or is removed, the MegaRAID firmware automatically rebuilds the redundancy of the virtual drive by providing an emergency hot spare (EHSP) drive, even if no commissionable dedicated drive or global hot spare drive is present.

4.8.9.1 Emergency Hotspare for Physical Drives

The Emergency Hotspare property determines whether a particular drive is capable of becoming a hot spare. This property appears under the controller properties only if the **Global Hotspare for Emergency**, and the **Unconfigured Good for Emergency** controller properties are enabled.

Follow these steps to view if the particular drive is capable of becoming an Emergency Hotspare.

1. Click the physical drive node in the right panel on the WebBIOS main dialog.

The Emergency Hotspare property of the drive is displayed, as shown in [Figure 116](#).

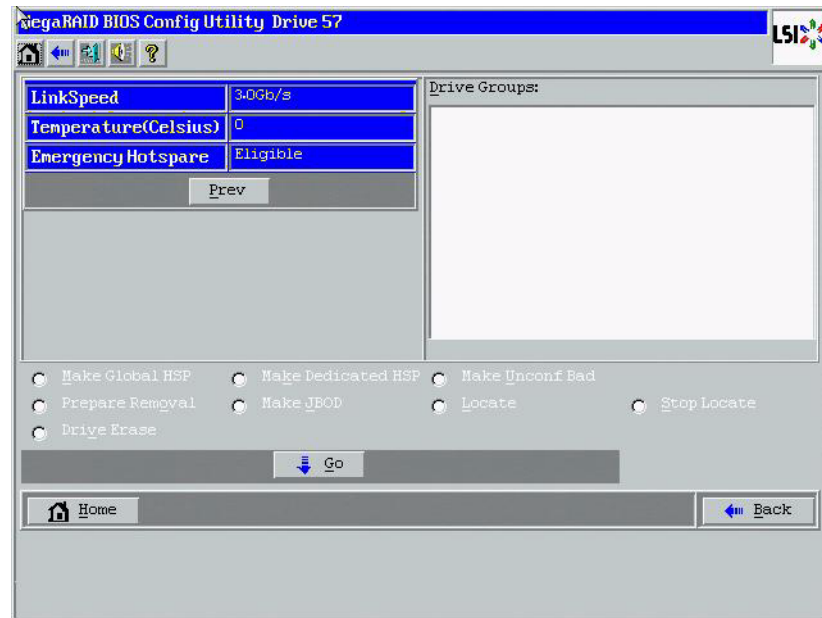


Figure 116: Emergency HotSpare

4.8.10 Emergency Hotspare for Controllers

The Emergency Hotspare properties are configured in the controller properties based on enabling or disabling the following properties:

- **Global Hotspare for Emergency**
- **Unconfigured Good for Emergency**
- **Emergency for SMARTer**

4.8.10.1 Setting Controller Hot Spare Properties

Follow these steps to set the Emergency Hotspare properties for controllers.

1. From the WebBIOS main menu, click **Controller Properties**.
2. Click **Next**.
3. Click **Next**.
4. Click **Next** again to go to the controller properties last property page and get all options to change the Emergency Hotspare properties.

The controller properties dialog appears, and you can enable or disable the properties, as shown in Figure 117.

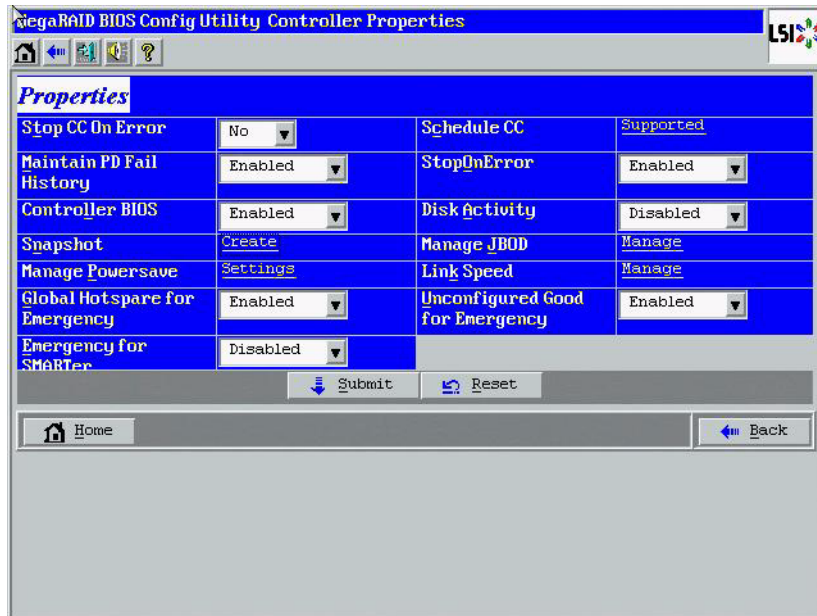


Figure 117: Setting Controller Hotspare Properties

4.8.10.2 Viewing Controller Hotspare Properties

Follow these steps to view the controllers Emergency Hotspare properties.

1. Click the **Controller properties** link in the WebBIOS main menu.
2. Click **Next**.

The Controller Properties are displayed as shown in [Figure 118](#).

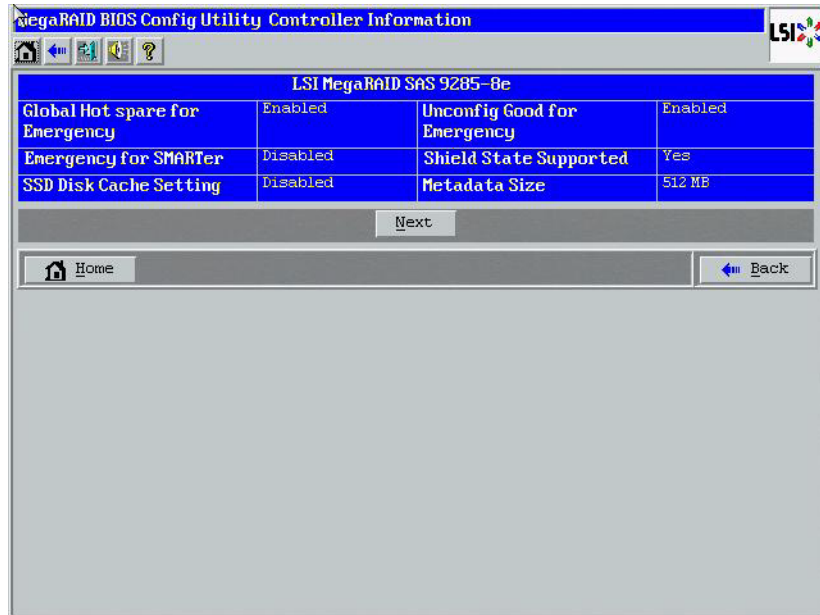


Figure 118: Viewing Controller Hotspare Properties

4.8.10.3 Commissioned Hotspare

The Commissioned Hotspare is used to determine whether the online drive has a Commissioned Hotspare drive assigned to it.

Click the online physical drive node in the right panel on the WebBIOS main dialog to view the Commissioned Hotspare property. See [Figure 119](#).

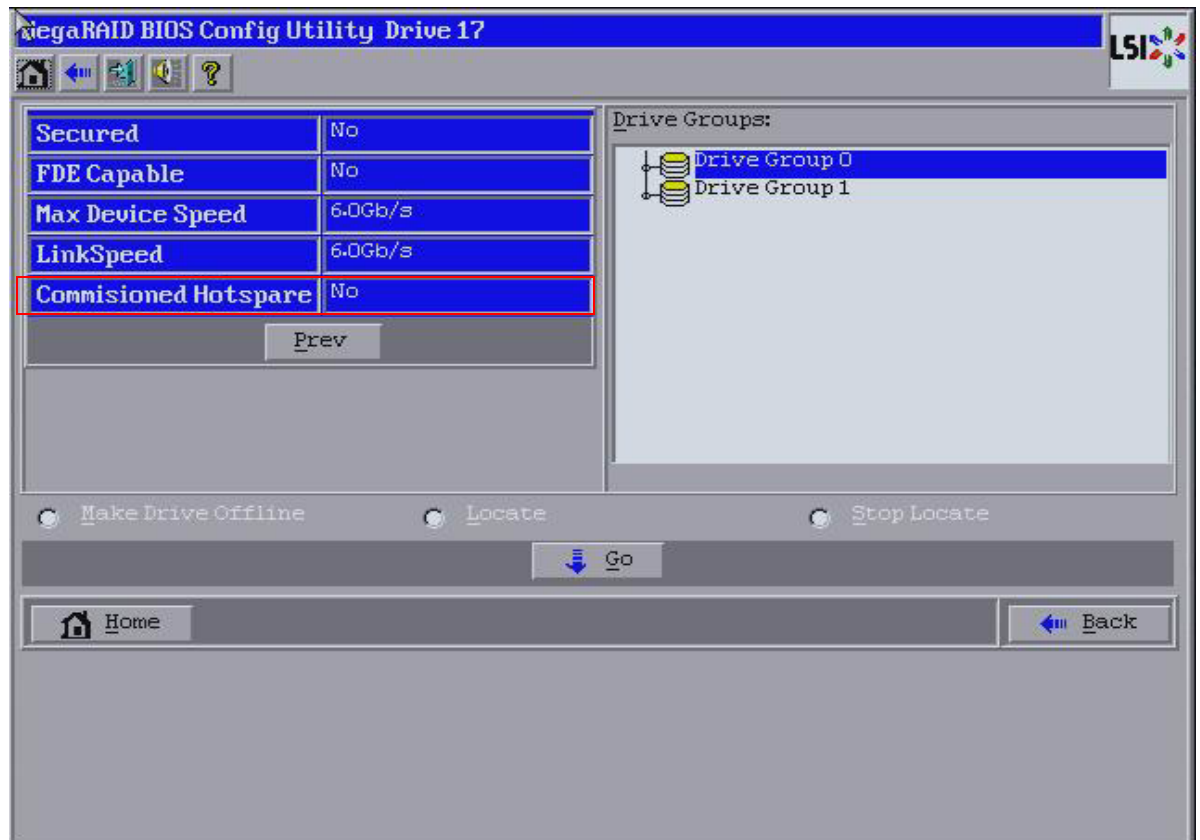


Figure 119: Commissioned Hotspare

4.9 Viewing and Expanding a Virtual Drive

Follow these steps to view virtual drive properties.

1. In the Logical view of the device tree, click the **Virtual Drive Node**.

The virtual drive properties are displayed, as shown in [Figure 120](#).

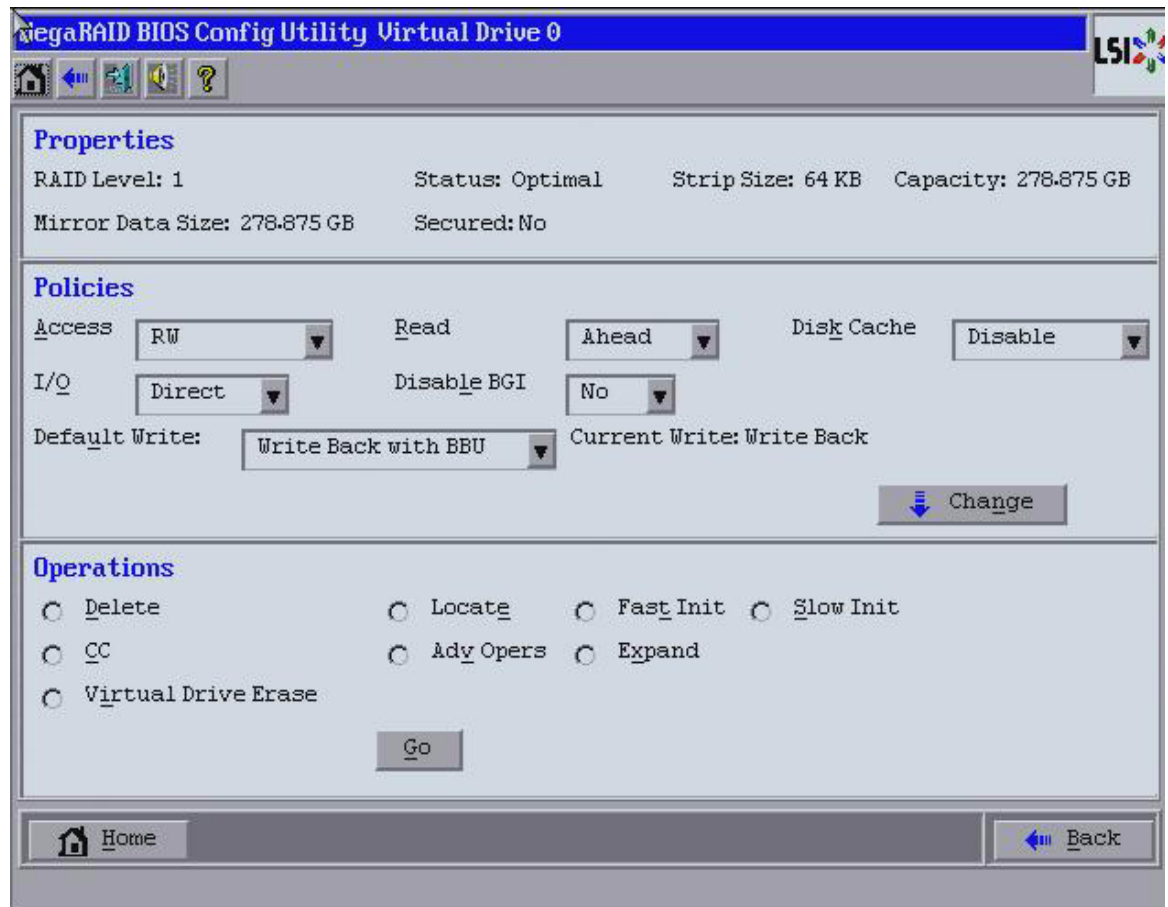


Figure 120: Virtual Drive Properties

You can increase the size of a virtual drive to occupy the remaining capacity in a drive group.

2. Select the **Expand** radio button, and click **Go**.

The WebBIOS Config Utility Expand Virtual Drive dialog appears, as shown in Figure 121.

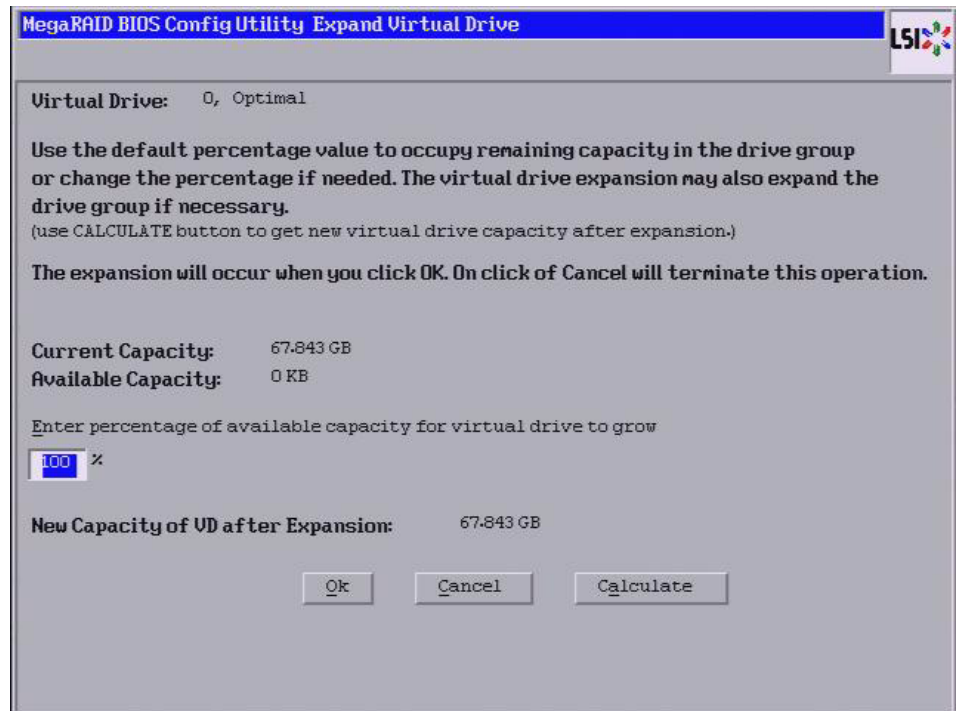


Figure 121: Expand Virtual Drive Dialog

3. Enter the percentage of the available capacity that you want the virtual drive to use.

For example, if there are 100 GB of capacity available and you want to increase the size of the virtual drive by 30 GB, select 30 percent.

4. Click **Calculate** to determine the capacity of the virtual drive after expansion.
5. Click **Ok**.

The virtual drive expands by the selected percentage of the available capacity.

4.10 Suspending and Resuming Virtual Drive Operations

MegaRAID provides background Suspend and Resume features that enhances the functionality. The background operations on a virtual drive can be suspended using the **Suspend** option, and later resumed using the **Resume** option. The suspended operation resumes from the point where the operation was suspended.

If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place it was stopped.

NOTE: Suspend and resume are applicable for all the background operations, such as background initialization, rebuild and consistency check notes.

Follow these steps to suspend an operation and resume an operation.

1. From the WebBIOS main menu, click **Virtual Drives** link.

or

From the task bar, click the **VD Progress Info** button.

The Virtual Drives main dialog appears, as shown in [Figure 122](#).

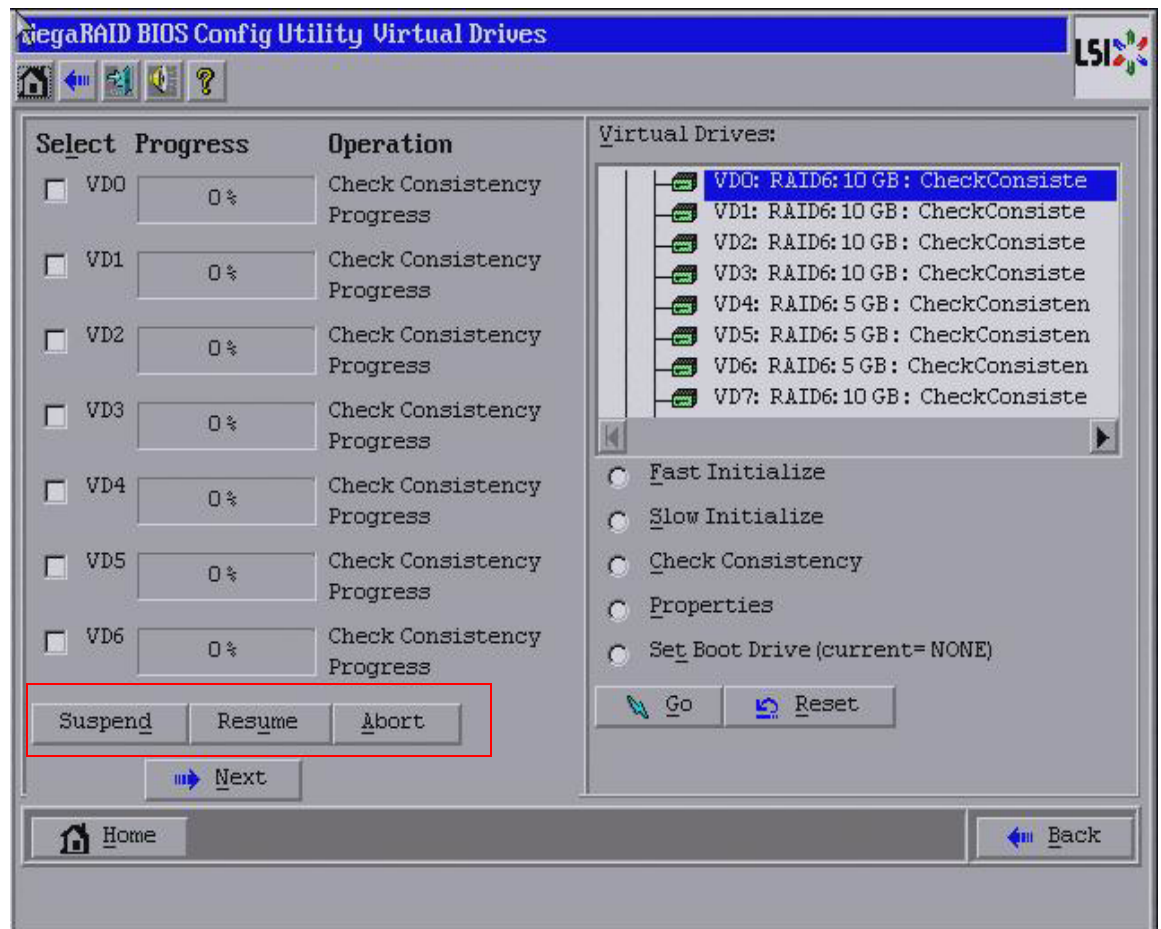


Figure 122: Virtual Drives Dialog

2. To suspend operations, select the check boxes for the operations that you want to suspend, and click **Suspend** (ALT+D).
3. To abort operations, select the check boxes for the operations and click **Abort** (ALT+A). Aborted operations cannot be resumed and have to be started again.
4. To resume operations, select the check boxes for the suspended operations that you want to resume, and click **Resume** (ALT+U).

4.11 Using MegaRAID Recovery

MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if data is deleted accidentally or maliciously, you can restore the data from the view or roll back to a snapshot at a previous point-in-time (PiT). MegaRAID Recovery supports up to eight snapshots of PiTs for each volume.

Each Recovery PiT volume snapshot is typically a fraction of the original volume size, because it tracks only the changes that are made to a volume after the PiT is created. Disk space for PiTs is reserved in the Snapshot Repository virtual drive, and the PiT is expanded in small increments as new data is written to the volume. Multiple PiTs of each volume can be retained online, enabling frequent snapshots to be stored in a space-efficient manner.

CAUTION: Do not select the virtual drive containing the operating system (OS) or any data as the snapshot repository. Updates to the operating system, operating system crashes, or any data updates could destroy data on that virtual drive.

4.11.1 Recovery Scenarios

Use the Recovery Features in three primary scenarios:

1. Restore the missing or deleted files (restore from view).
 - a. Discover the files are missing or deleted.
 - b. Review the snapshot views of the file content (also known as mounting a snapshot) from each PiT until you find the missing file.

A snapshot view contains the content from the point-in-time at which the snapshot was made.
 - c. Drag and drop the missing file from snapshot view back into the online storage volume that was the source of the snapshot.
2. If corrupt operating system files exist in a volume, roll back the volume to a previous state.
 - a. Reboot the system, and run WebBIOS.
 - b. Select the most recent snapshot that does not contain the corrupted or malicious file to roll back to. Select the most recent PiT snapshot to roll back to.
 - c. Reboot the system.

The system automatically rolls back to its previous state based on the selected PiT snapshot.
3. Reduce the risk of extended downtime during application updates/upgrades in the IT center.
 - a. When the application is offline, take a snapshot of the application volume.
 - b. Install each patch individually, and test for any new defects that might have been introduced.
 - c. Take a snapshot after you test each patch, and determine that it is clean.
 - d. If a defect is introduced, roll back to the previous installation, and bypass the installation of the defective patch.

NOTE: If the volume is still damaged, continue to select from the next most current PiT snapshot to the oldest.

4.11.2 Enabling the Recovery Advanced Software

You can enable the Recovery advanced software in WebBIOS. After you enable Recovery, you create two virtual drives—one as a snapshot base or source and the other as a snapshot repository. The snapshot base virtual drive contains the data that is stored in the snapshot repository virtual drive.

Follow these steps to enable MegaRAID Recovery.

1. Click a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog to access the Virtual Drive dialog.

The Virtual Drive dialog appears, as shown in [Figure 123](#).

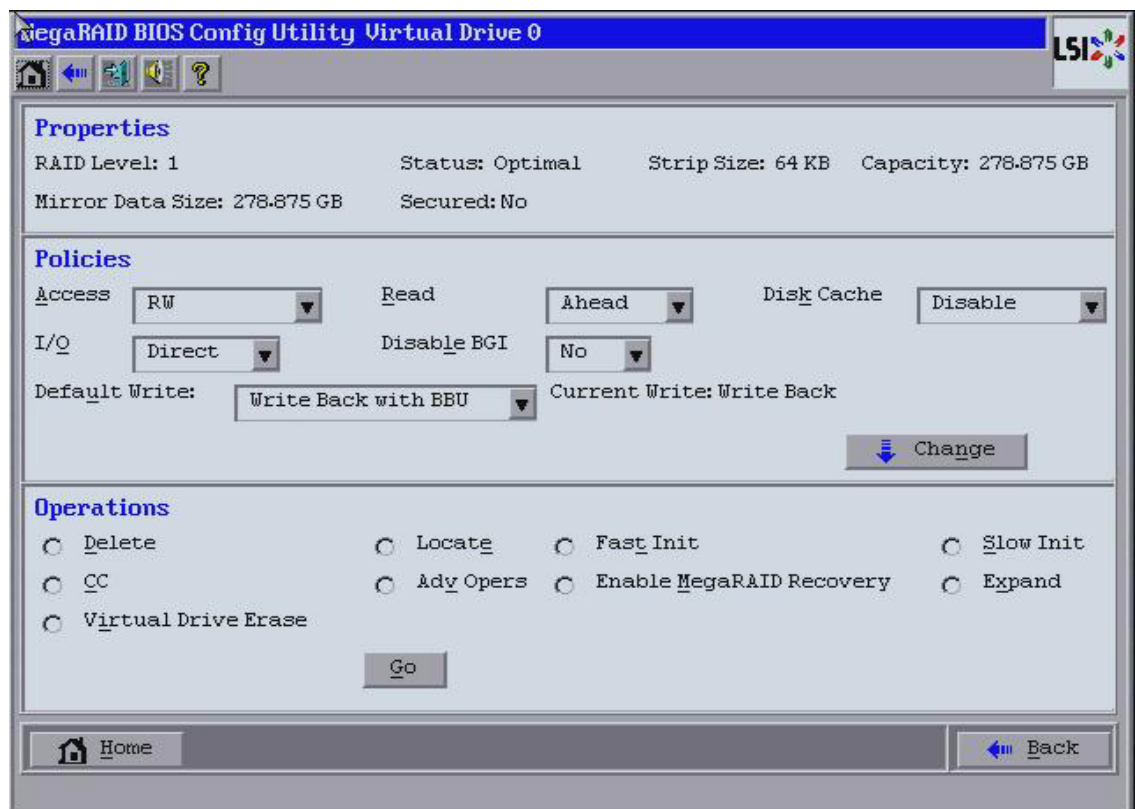


Figure 123: Virtual Drive Dialog

2. Click **Enable MegaRAID Recovery** in the Operations panel of the dialog.

3. Click **Go** in the Operations panel of the dialog.

The Enable MegaRAID Recovery dialog appears, as shown in [Figure 124](#).

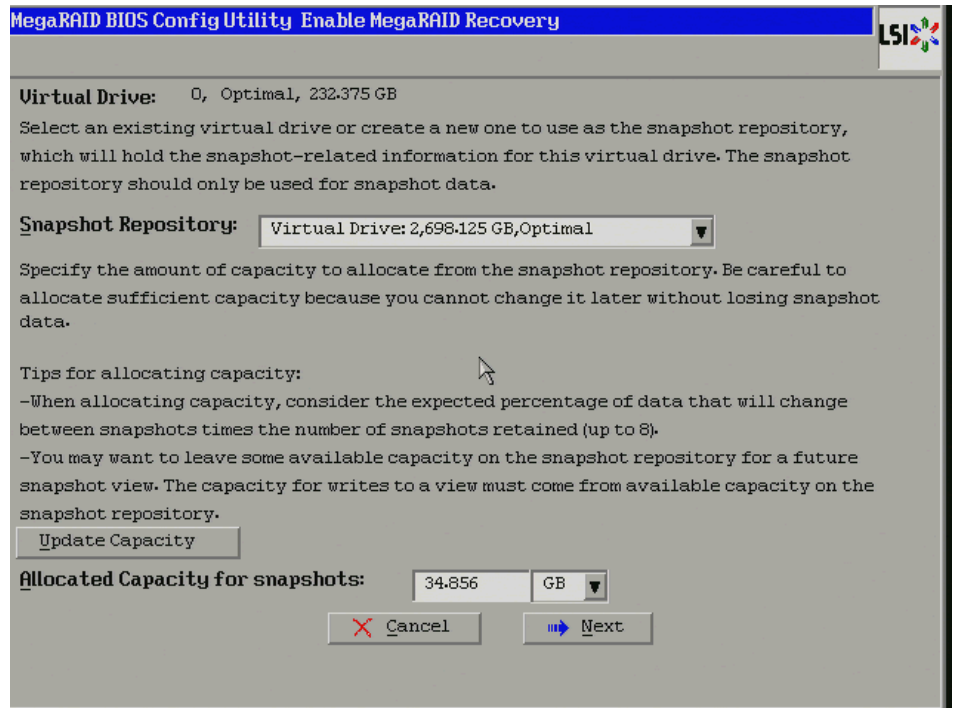


Figure 124: Enable MegaRAID Recovery Dialog

4. Select a virtual drive from the list of virtual drives in the **Snapshot Repository** drop-down list.

This setting is the snapshot repository virtual drive. This drive stores the snapshot data. Make sure you select a snapshot repository virtual drive with enough available capacity. The available capacity is the largest free block of capacity on the selected repository.

NOTE: A virtual drive and a snapshot repository virtual drive can be associated with the same drives or a common set of drives, or the two virtual drives can be located on two completely separate set of drives. Using a separate set of drives for the virtual drive and the snapshot repository virtual drives provides a performance advantage over using a common set of drives.

5. Click the **Update Capacity** button to determine the available capacity of the selected repository.

CAUTION: Do not select the virtual drive containing the operating system as the snapshot repository. Updates to the operating system crashes can destroy data on that virtual drive.

6. In the **Allocated Capacity for snapshots** field, select the available capacity in the snapshot repository to use for changes to the virtual drive.

The capacity is dependent on how write-intensive the application is of which you are taking snapshots. The available capacity is the largest free block of capacity on the snapshot repository virtual drive.

NOTE: If you use all of the space of the snapshot repository virtual drive, you cannot space to create a snapshot and a view, because of insufficient space.

CAUTION: Copy all of your data to another virtual drive before you select this option. If any existing data exists on this virtual drive, it will be lost.

7. Click **Next**.

The snapshot settings dialog appears, as shown in [Figure 132](#).

8. Click **Finish**.

A confirmation dialog appears.

9. Confirm that you want to make these selections.

This virtual drive becomes a snapshot repository. Use it only for storing snapshot-related data.

CAUTION: After you enable snapshots on this virtual drive, you cannot change the allocated percentage of capacity or the snapshot repository without first disabling snapshots and losing any snapshot data.

4.11.3 Creating Snapshots and Views

You can use WebBIOS to create up to eight snapshots of a volume. WebBIOS shows the snapshots in chronological order from the oldest to the newest. Each snapshot is a PiT snapshot of the virtual drive that is the snapshot base. First, create the snapshot base virtual drive, and then create the snapshot.

After you create the snapshots, you can create views of the PiT snapshots. You can search the views to find a snapshot that does not contain the corrupt data or a snapshot that contains the deleted data, depending on the situation. After you create a snapshot, you can reboot and roll back to a snapshot to restore data.

Follow these steps to create a snapshot.

1. Enable MegaRAID Recovery.

See [Section 4.11.2, *Enabling the Recovery Advanced Software*](#), for the procedure used to enable MegaRAID Recovery in WebBIOS.

2. Click on the virtual drive in the Logical View on the main dialog to go to the operations for the virtual drive.

The Virtual Drive dialog appears, as shown in Figure 125.

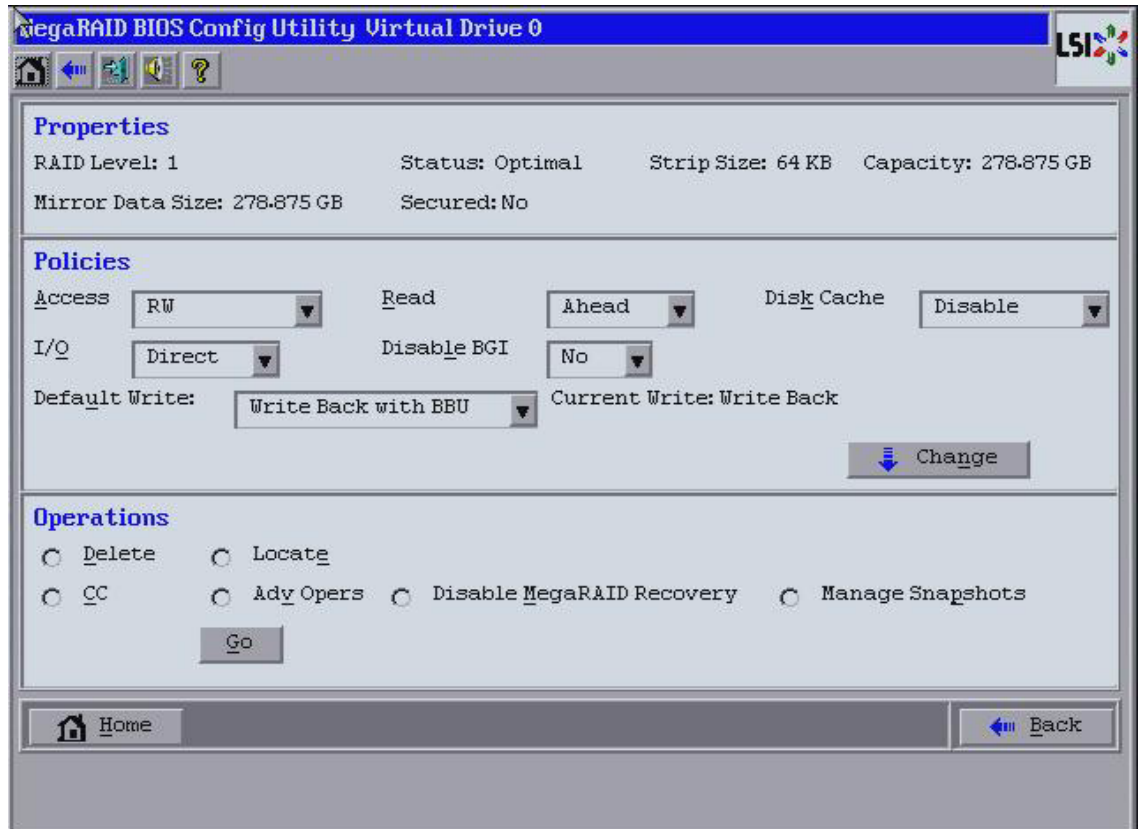


Figure 125: Snapshot Base Virtual Drive Dialog

3. Click **Manage Snapshots** in the Operations panel.

The Virtual Drive Properties dialog appears, as shown in [Figure 126](#).

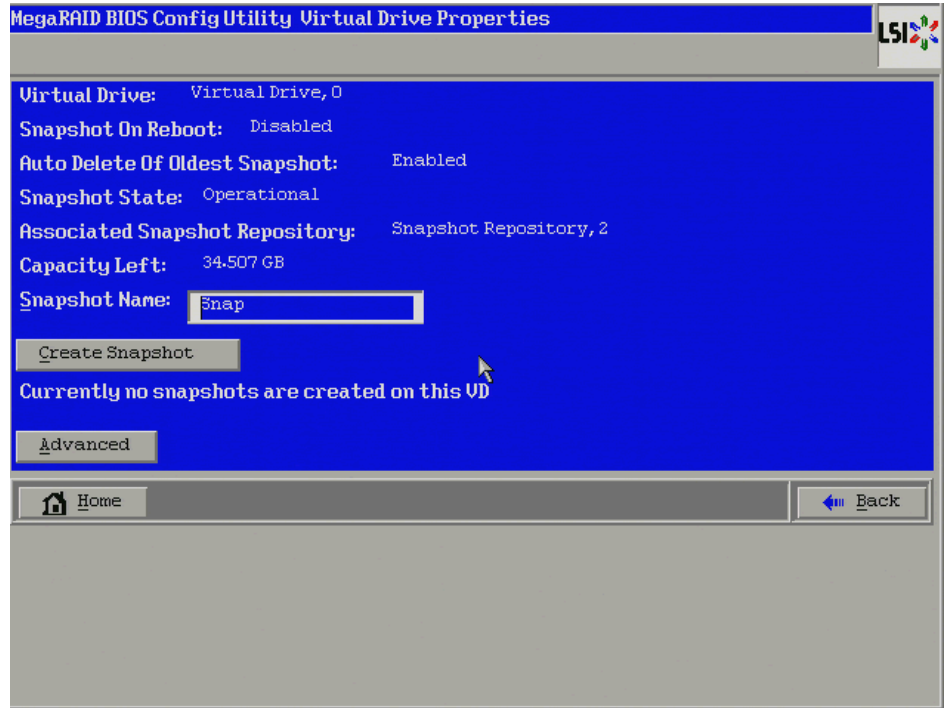


Figure 126: Virtual Drive Properties Dialog

4. Enter a snapshot name in the Snapshot name field and click **Create Snapshot**.

This action creates a snapshot that appears as a link in the **Snapshot Timeline**.

5. Click the link of a specific snapshot.

The snapshot details appear.

6. Click **Advanced**.

The Snapshot Settings dialog appears, as shown in [Figure 132](#).

7. Click **Create View**.

The Create View dialog appears, as shown in [Figure 127](#).

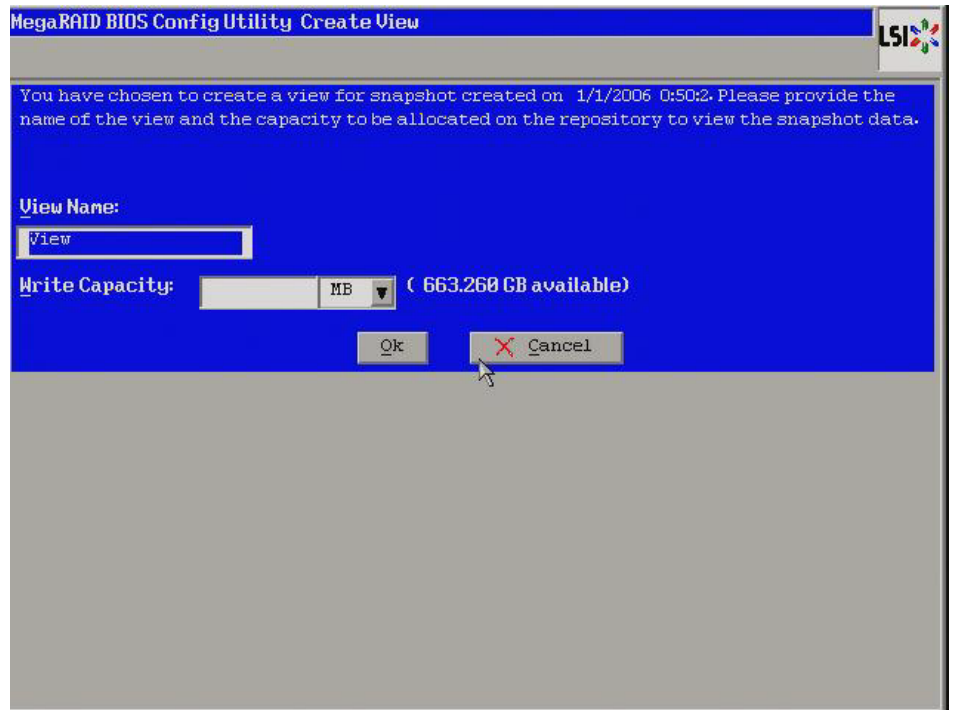


Figure 127: Create View Dialog

8. Enter a view name in the **View Name** field, specify the capacity of the view in the **Write Capacity** field and click **Ok**.

This action creates the view. After you create a view, you can view details about both the snapshot and the view on a single page, as shown in [Figure 128](#).

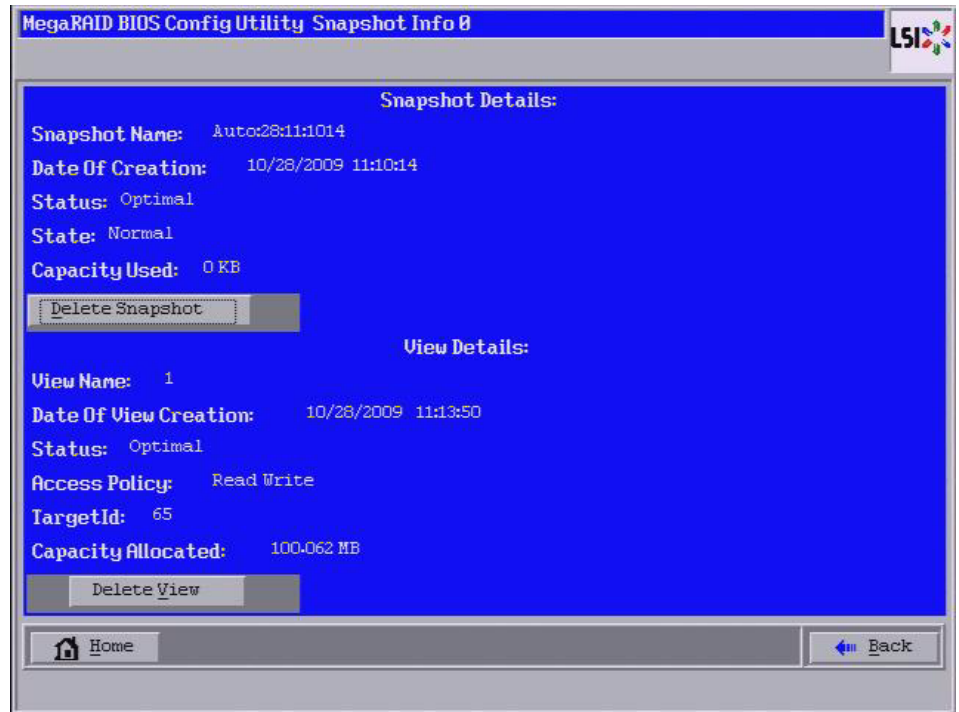


Figure 128: Snapshot Info Dialog

4.11.4 Creating Concurrent Snapshots

If you have created multiple snapshot base virtual drives, you can create snapshots on all of them at one time (concurrent snapshots). Each snapshot has the same name and time stamp.

Follow these steps to create concurrent snapshots.

1. Click **Controller Properties** on the WebBIOS main dialog.

The first Controller Properties dialog appears. There are four Controller Properties dialogs.

- Click **Next** to on the first three Controller Properties dialogs to access the fourth Controller Properties dialog, as shown in Figure 129.

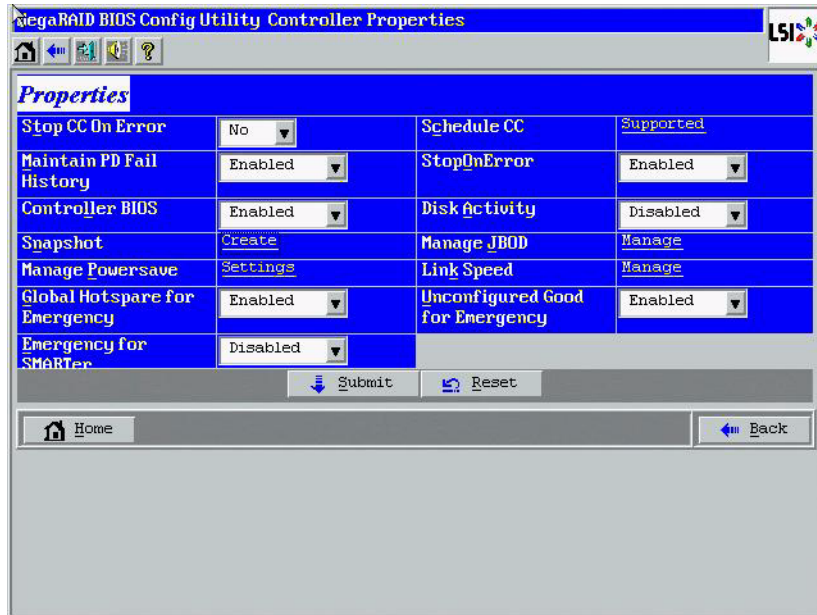


Figure 129: Fourth Controller Properties Dialog

- Click **Create** in the **Snapshot** field.

The Create Snapshots dialog appears, as shown in Figure 130.

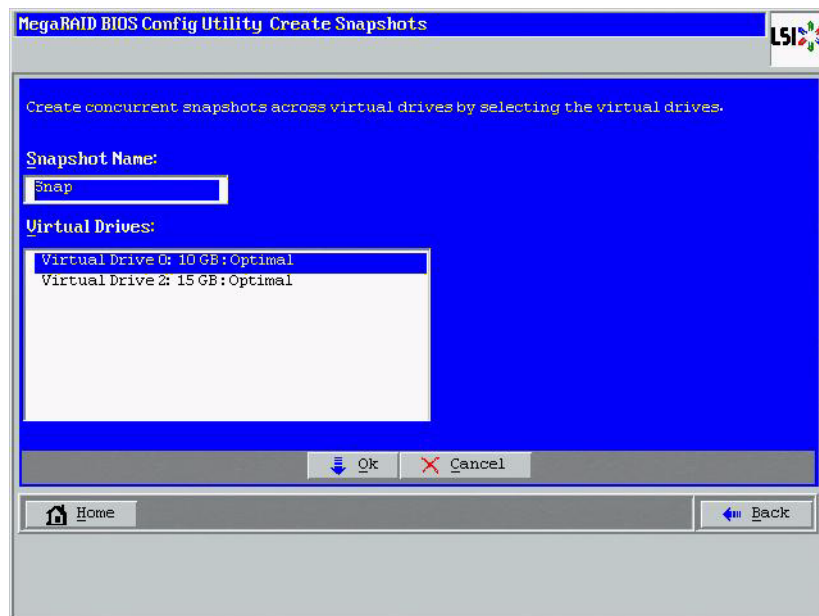


Figure 130: Create Snapshots Dialog

- Enter a snapshot name in the **Snapshot Name** field.

5. Select the virtual drives on which you want to create concurrent snapshots.
6. Click **OK**.

This action creates a snapshot with same name and the same timestamp on all of the selected snapshot base virtual drives.

4.11.5 Selecting the Snapshot Settings

You can use the Snapshot Settings dialog to perform the following actions:

- Take a snapshot on reboot.
This action takes a snapshot of the virtual drive when you reboot after every successful system shutdown. This feature is mainly intended to take a snapshot of boot virtual drives to allow the operating system to be restored in case of corruption.
- Enable automatic deletion of a snapshot.
This action deletes the oldest snapshot automatically and lets you create a new snapshot.

Follow these steps to enable the snapshot settings.

1. Click a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.

The Virtual Drive dialog appears, as shown in [Figure 131](#).

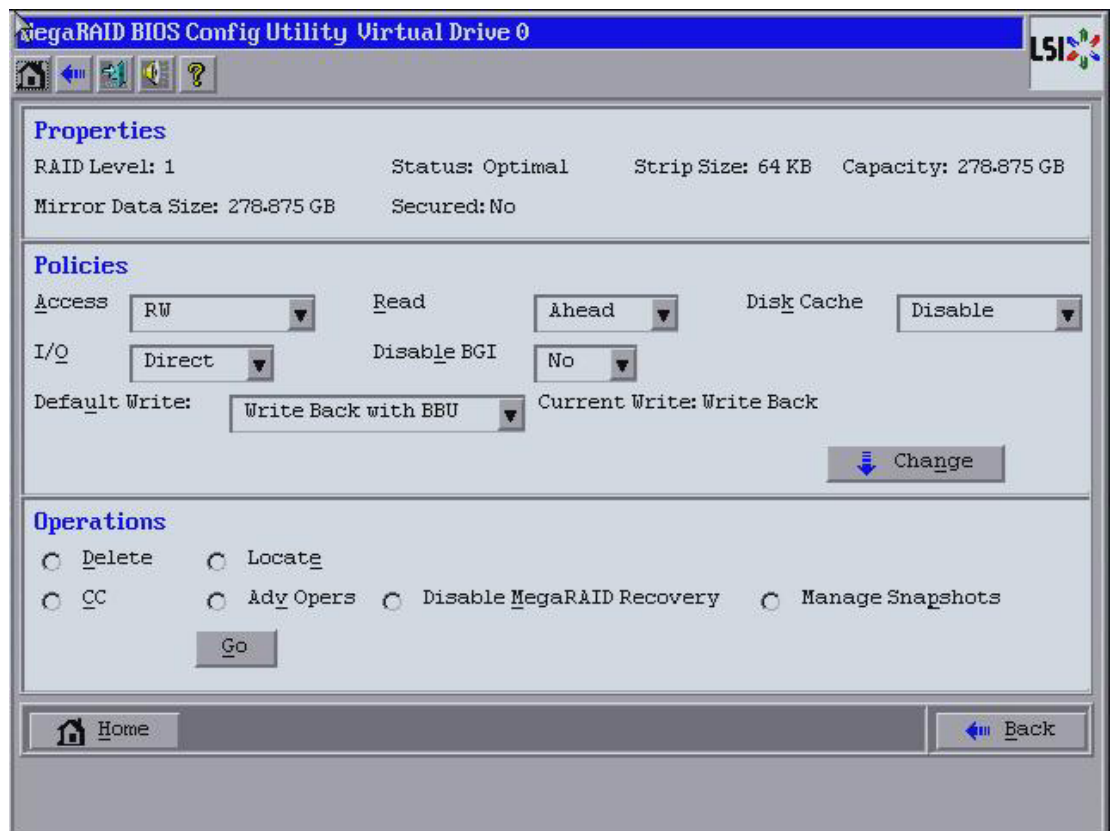


Figure 131: Virtual Drive Dialog

2. Click **Manage Snapshots**.

The Snapshot Settings dialog appears, as shown in [Figure 132](#).

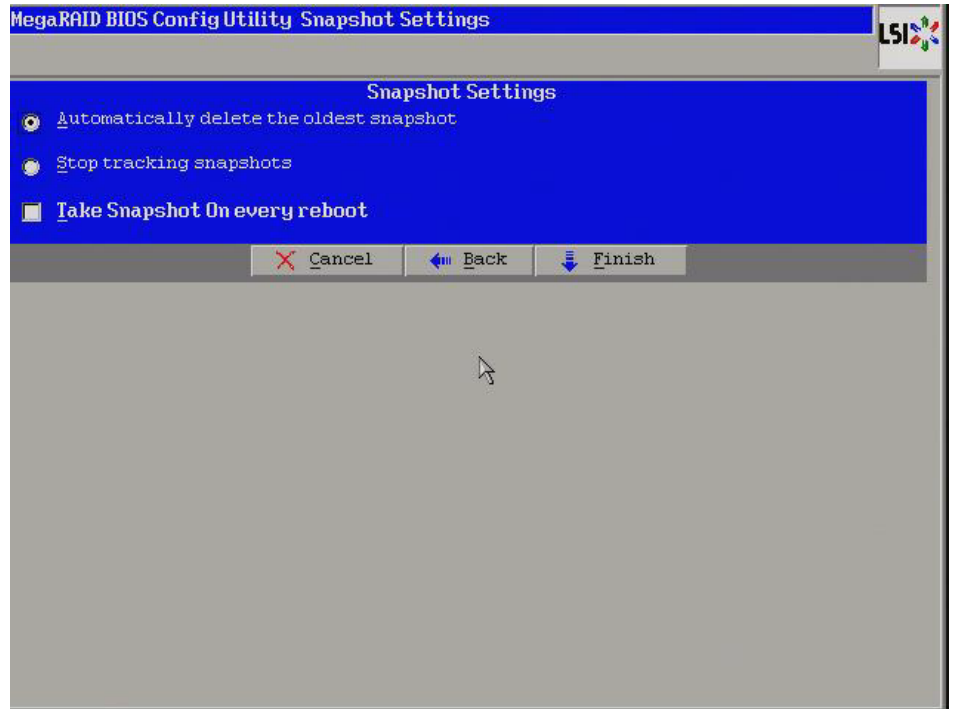


Figure 132: Snapshot Settings Dialog

3. Select the **Take Snapshot on every Reboot** check box, or select the **Automatically delete the oldest snapshot** radio button, or select the **Stop tracking snapshots** radio button.
4. Click **Finish**.

4.11.6 Viewing Snapshot Properties

You can view the properties of a snapshot, such as the total capacity, capacity used, and capacity available.

Follow these steps to view snapshot properties.

1. Click a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.

The Virtual Drive dialog appears, as shown in [Figure 133](#).

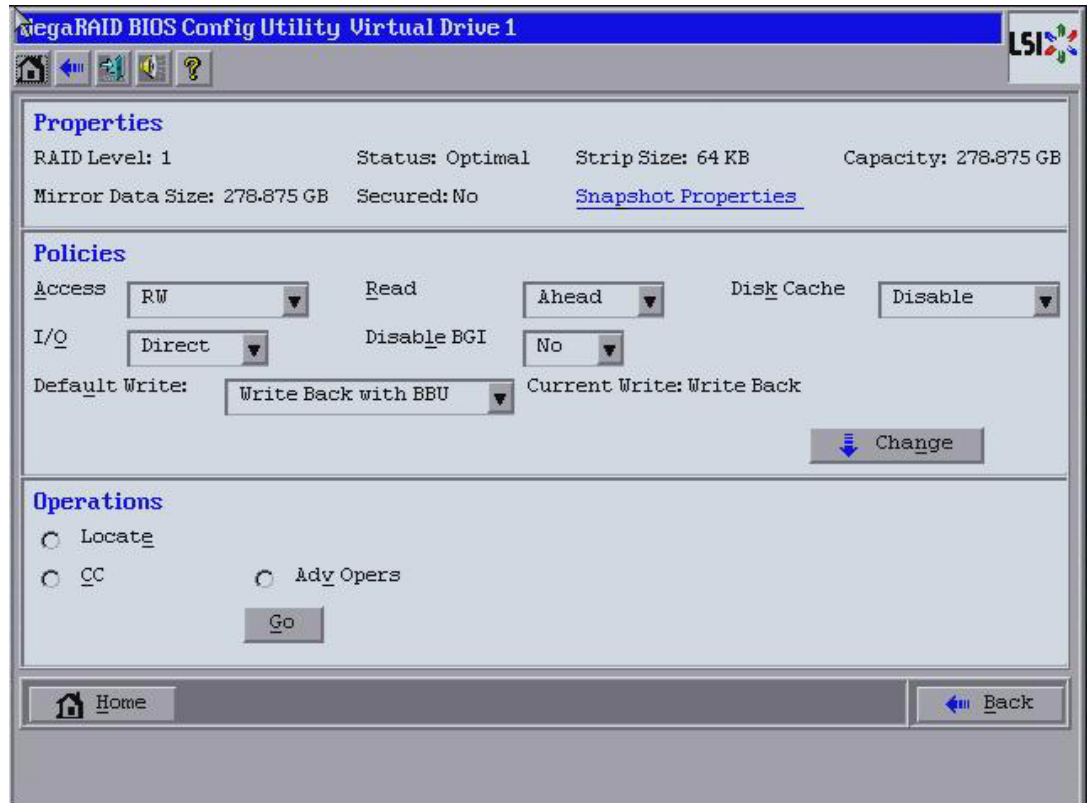


Figure 133: Virtual Drive Dialog

2. Click **Snapshot Properties**.

The Snapshot Repository Properties dialog appears, as shown in [Figure 134](#).

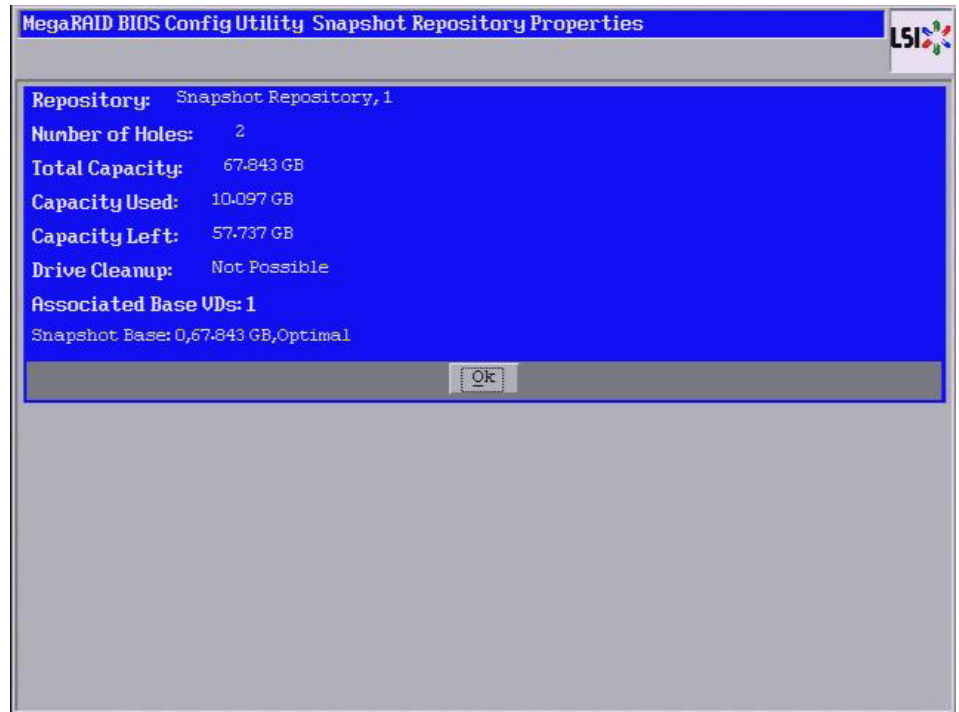


Figure 134: Snapshot Repository Properties Dialog

3. Click **OK** to return to the Virtual Drive dialog.

4.11.7 Restoring a Virtual Drive by Rolling Back to a Snapshot

You can roll back to a previous point-in-time snapshot to recover an entire volume. This action is often used where there are malicious files that cannot be traced. Reboot the system, and then roll back to a snapshot that does not have the malicious or corrupt files.

Follow these steps to roll back the volume version to an earlier version.

1. After you determine there are malicious or corrupt files, start the WebBIOS configuration utility.
2. Access the Virtual Drive dialog by clicking on a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.

The Virtual Drive dialog appears, as shown in [Figure 135](#).

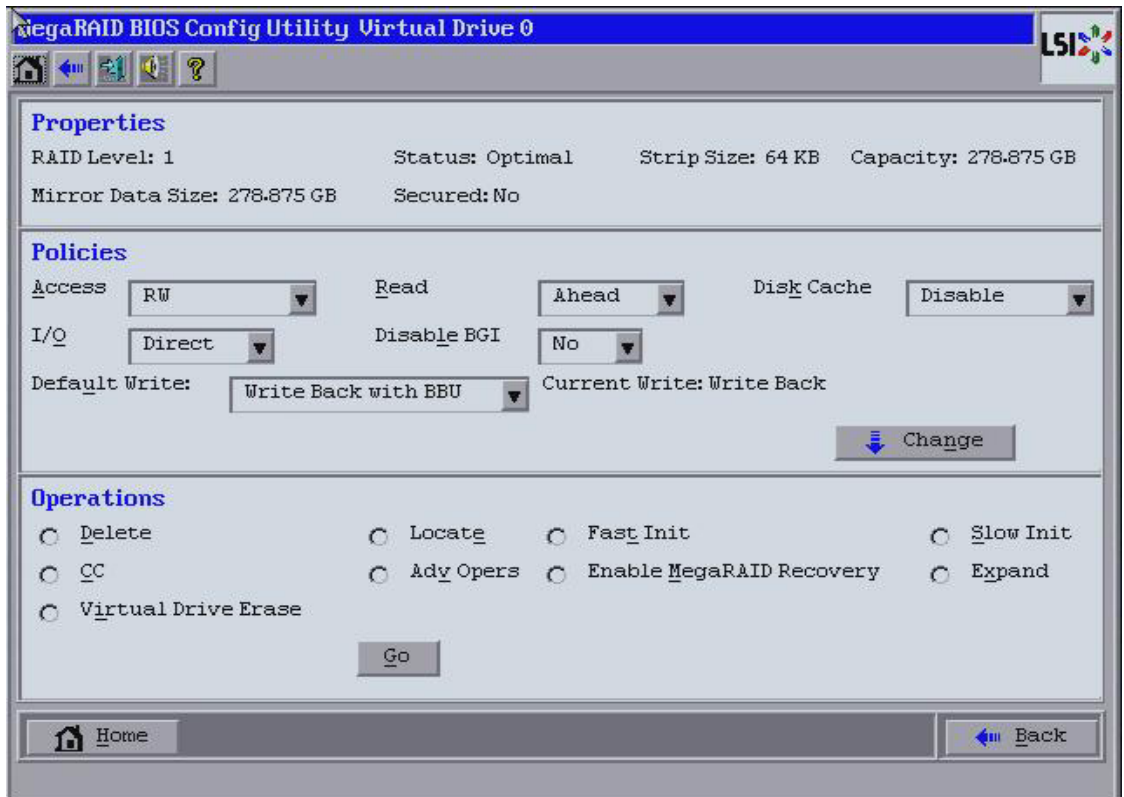


Figure 135: Virtual Drive Dialog

3. Click the **Adv Opers** radio button and then click **Go**.

The Advanced Operations dialog appears, as shown in [Figure 136](#).

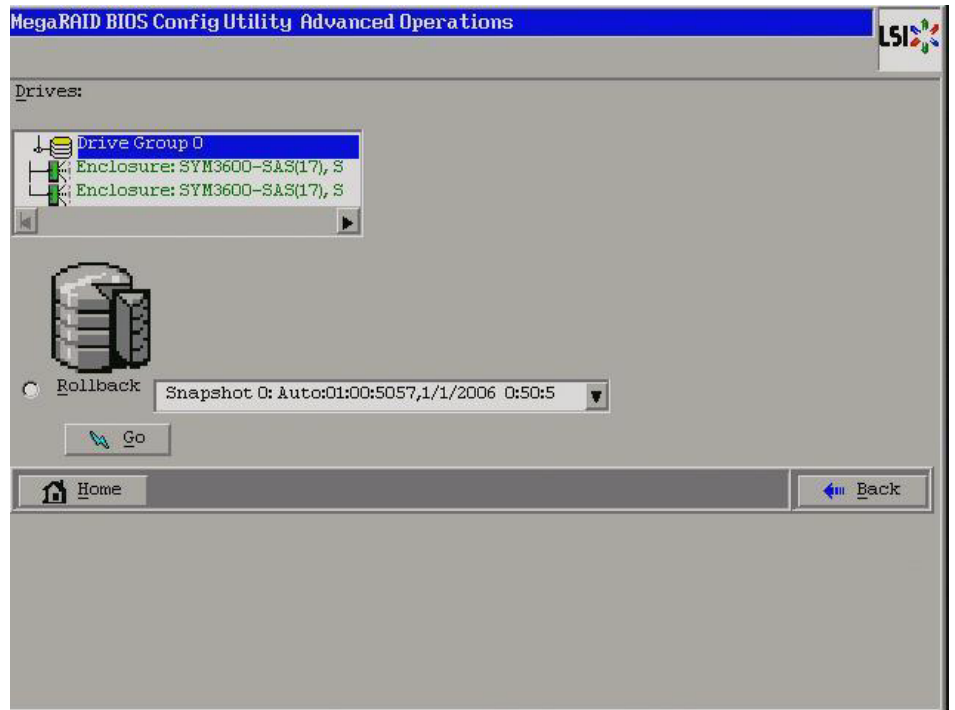


Figure 136: Advanced Operations Dialog

4. Select a snapshot from the drop-down list.

If the volume is still damaged, continue to select from the next most current PiT snapshot to the oldest.

5. Click **Go**.

The system rolls back to the selected PiT snapshot and returns you to a snapshot that does not have the malicious or corrupt files.

4.11.8 Cleaning Up a Snapshot Repository

The clean up option can be performed only on a snapshot repository virtual drive. Perform a cleanup if a snapshot base virtual drive goes offline and the snapshot repository virtual drive is still connected to the system. After you perform the cleanup, memory that was allocated to the offline base virtual drives will be available to the snapshot repository virtual drive.

Follow these steps to clean up a snapshot repository.

1. Access the Virtual Drive dialog by clicking a snapshot repository virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.

The Virtual Drive dialog appears, as shown in [Figure 137](#).

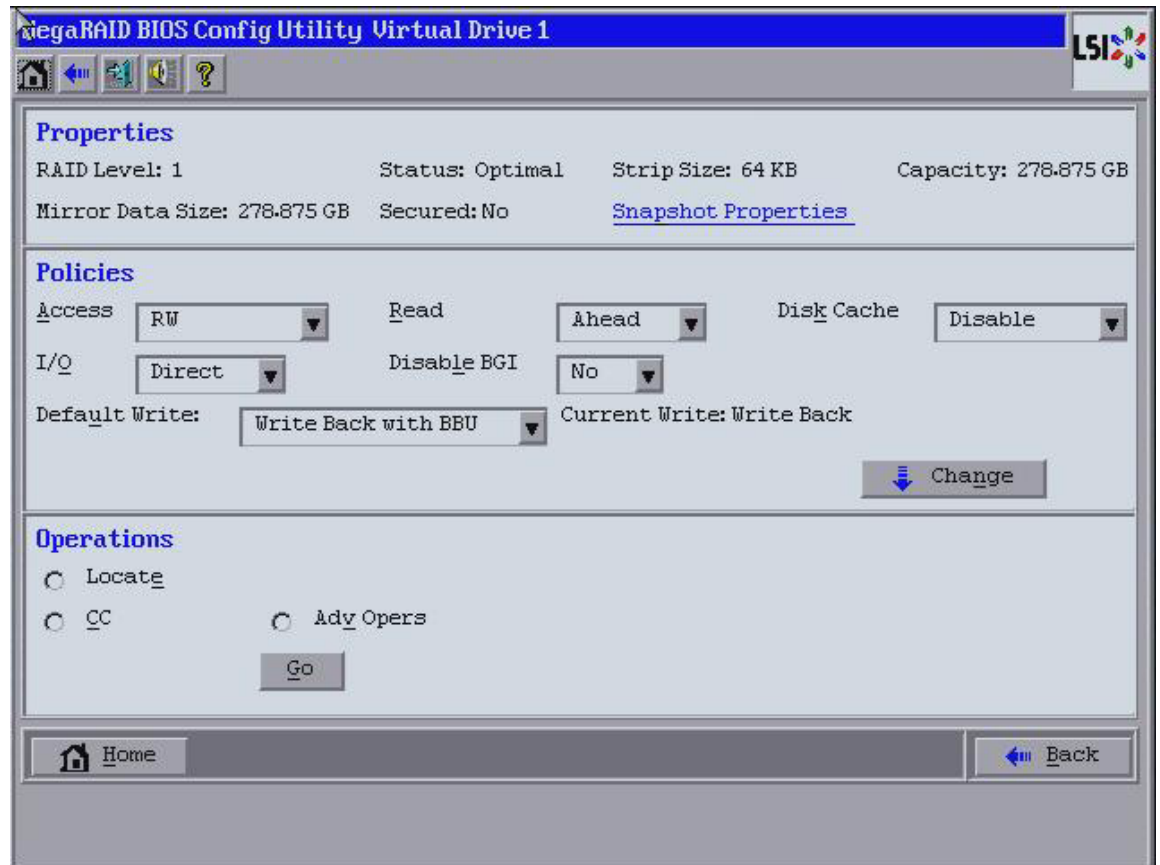


Figure 137: Virtual Drive Dialog

2. Click the **AdvOpers** radio button and then click **Go**.

The Advanced Operations dialog appears, as shown in [Figure 138](#).

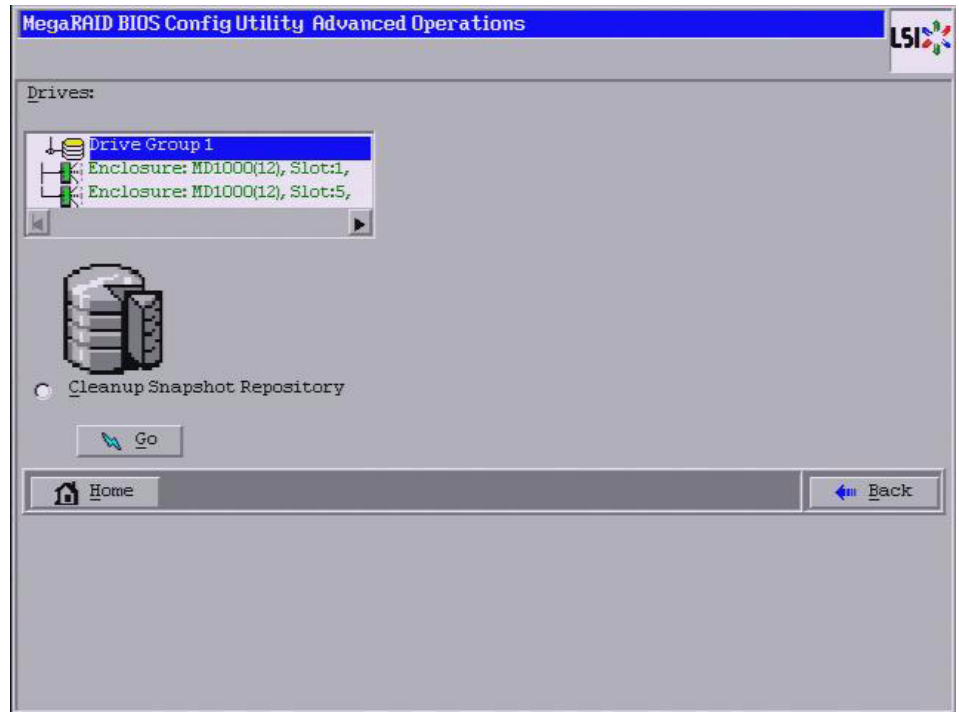


Figure 138: Advanced Operations Dialog

3. Select the **Cleanup Snapshot Repository**.
4. Click **Go**.

This action cleans up the snapshot repository.

4.12 Non-SED Secure Erase

4.12.1 Erasing a Non - SED Physical Drive

This section describes the procedure used to securely erase data on non self-encrypting drives (Non-SED), which are normal HDDs.

Follow these steps for non-SED secure erase.

1. Go to the Physical view in the WebBIOS main menu.
2. Click the physical drive node.

3. Select the **Drive Erase** radio button, as shown in [Figure 139](#) and click on **Go**.

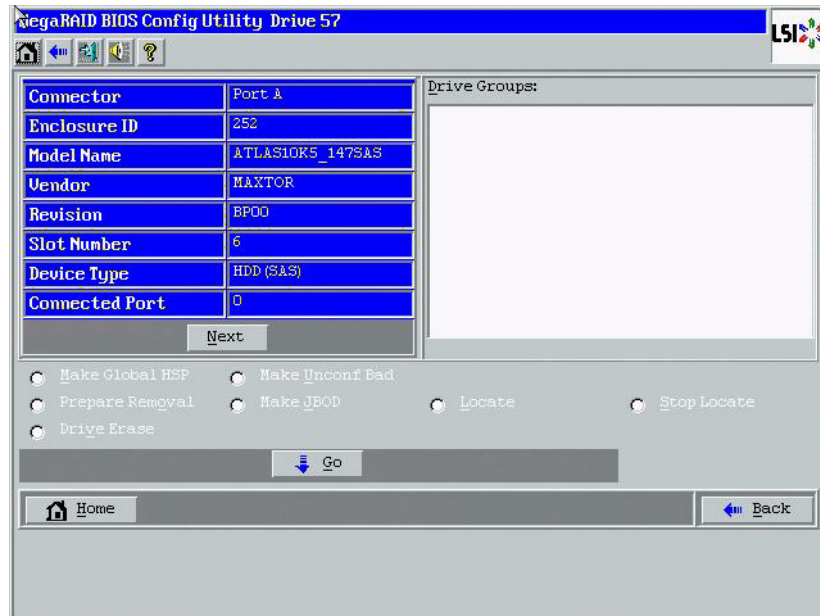


Figure 139: Physical Drive Dialog

The Mode Selection- Drive Erase dialog appears.

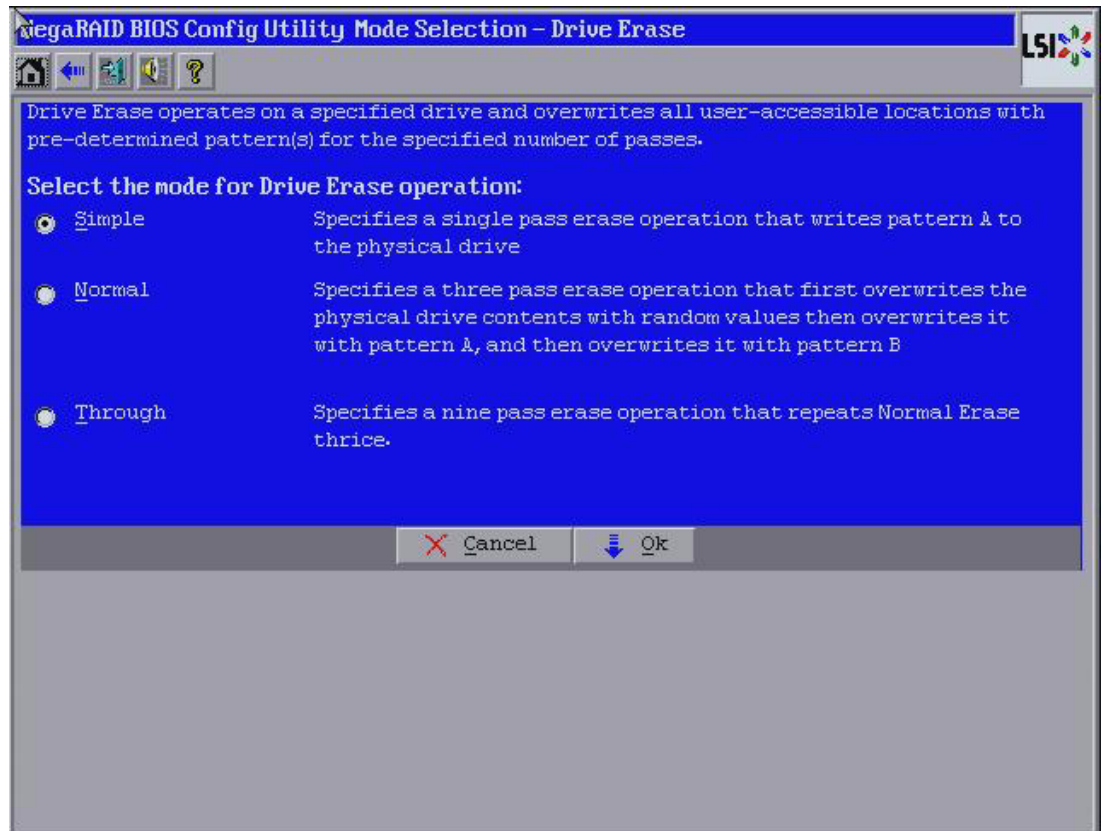


Figure 140: Mode Selection - Drive Erase

4. Select any of the modes available under the **Select the mode for Drive Erase Operation**
 - **Simple** – (Alt + S)
 - **Normal** – (Alt + N)
 - **Thorough** – (Alt + T)
5. Click **OK**. A confirmation message dialog appears as shown below.

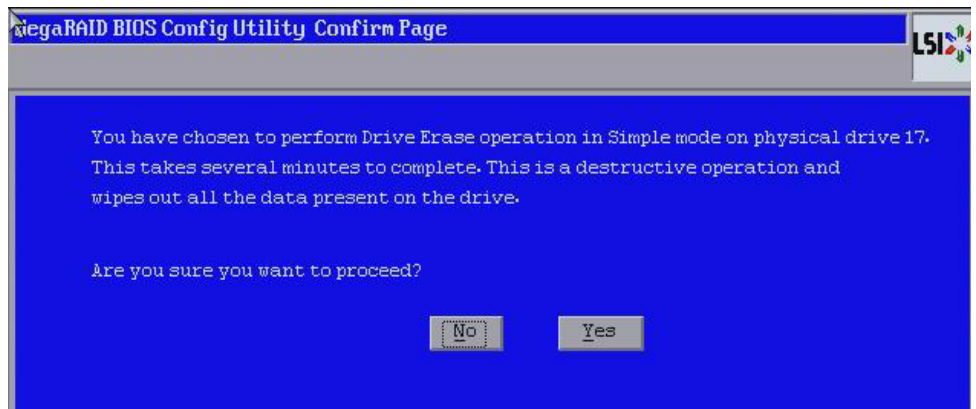


Figure 141: Drive Erase Confirm Page

4.12.1.1 Drive Erase Progress

Physical drives, erase operation is generally a time-consuming operation and is performed as a background task.

Follow these steps to check the progress of a physical drive erase operation.

1. Click the **Drives** link in the left panel on the WebBIOS main dialog.

The **Drive Erase Progress** appears, as shown in Figure 142.

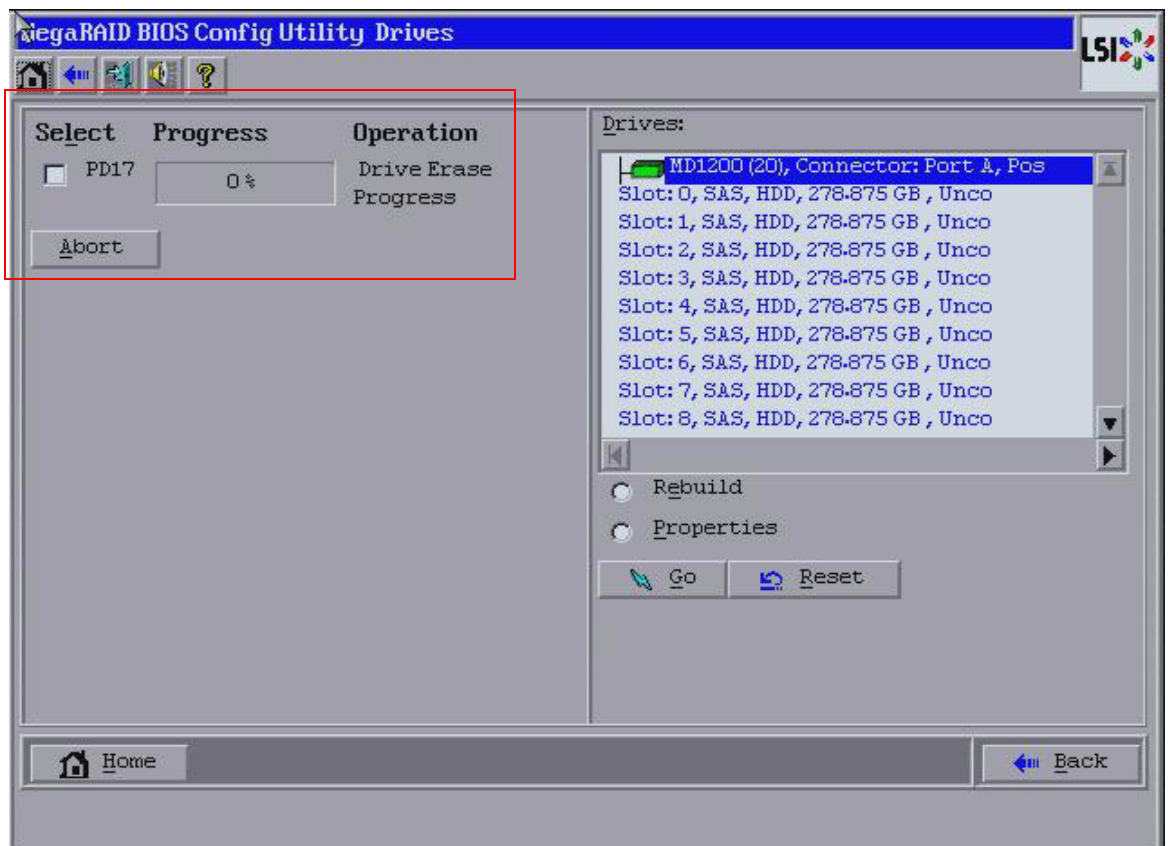


Figure 142: Drive Erase Progress

- To abort drive erase, select the check box for the operation that you want to abort and click **Abort**.

4.12.2 Virtual Drive Erase

Virtual drive erase is a background operation.

Follow these steps to perform the **Virtual Drive Erase** operation.

- Go to the **Logical view**.
- Click on the Virtual Drive node.
- Select the **Virtual Drive Erase** radio button and click **Go** as shown in [Figure 143](#).

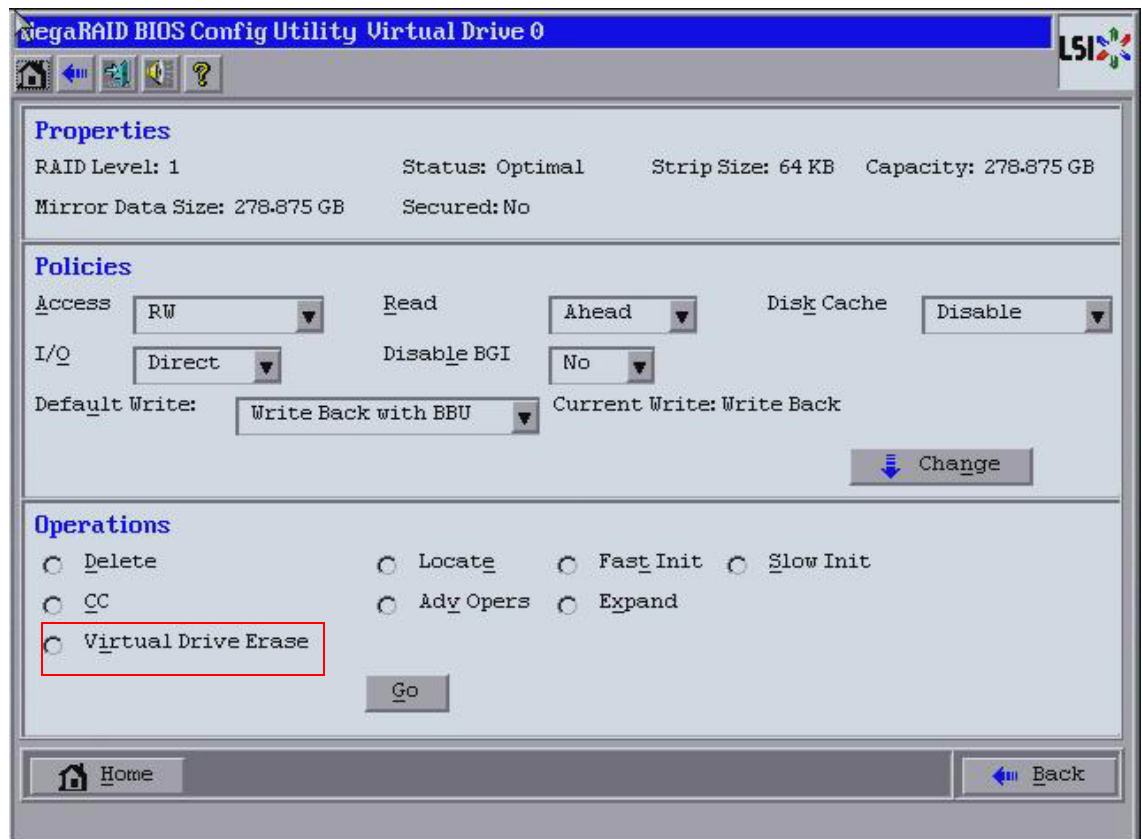


Figure 143: Virtual Drive Dialog

NOTE: The radio button appears for the unsecured virtual drives that are online.

The Mode Selection - Drive Erase Dialog is displayed, as shown in Figure 144.

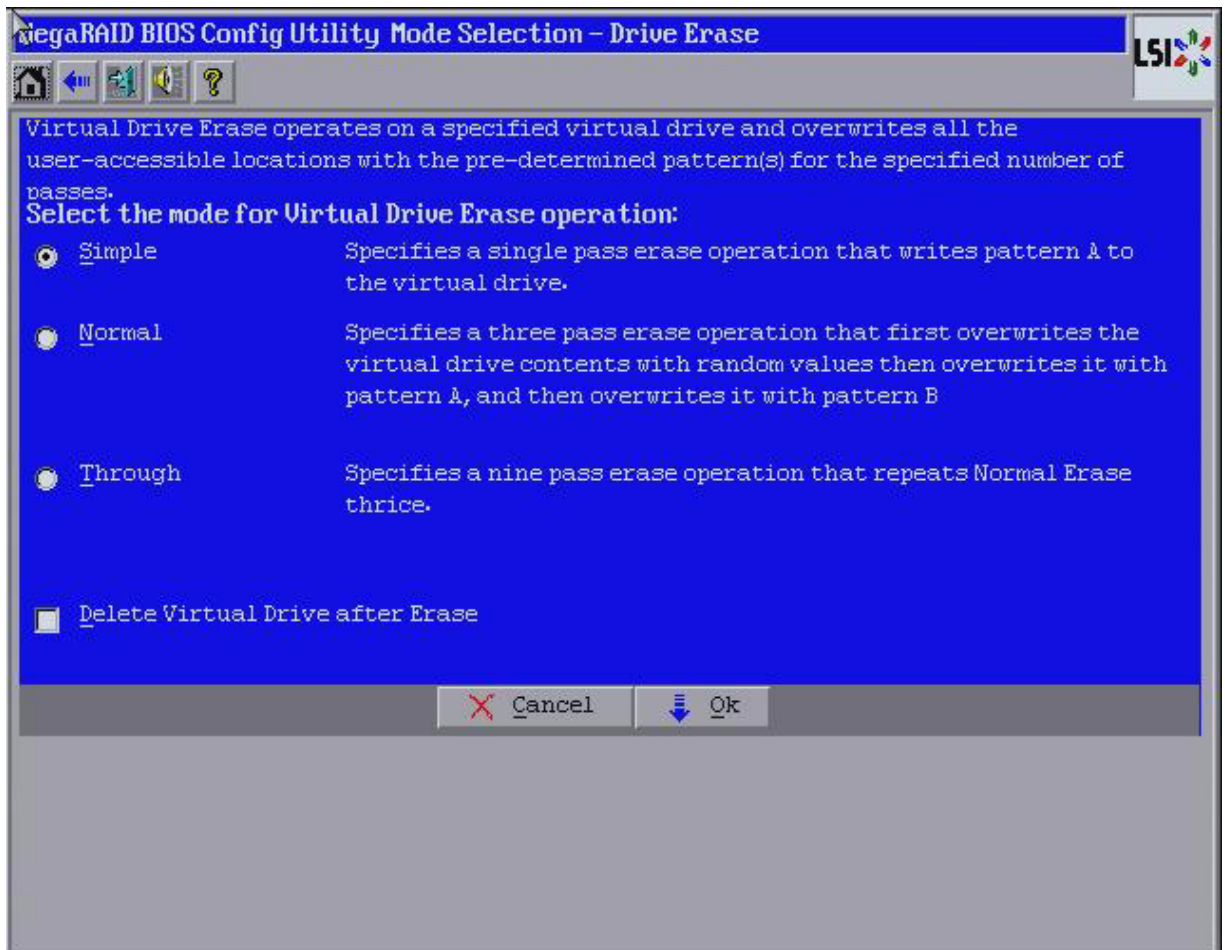


Figure 144: Virtual Drive Erase Dialog

4. Select any of the following options.
 - **Simple** – (**Alt + S**) –After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Normal** – (**Alt + N**) – After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Thorough** – (**Alt + T**) –After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Delete Virtual Drive after Erase** (**Alt + D**) - if you select this check box, the virtual drive is erased, and a confirmation dialog appears.
 - **OK** (**Alt + O**) – Click **OK** and, if the **Delete Virtual Drive after Erase** check box is selected a confirmation dialog appears.
 - **Cancel** (**Alt + C**) – Clicking this option, closes the dialog, and the WebBIOS navigates back to the Virtual Drive dialog.

4.12.2.1 Group Show Progress for Virtual Drive Erase

The virtual drive erase operation is a time-consuming operation, and it is performed as a background task.

Follow these steps to view the progress of **Virtual Drive Erase**.

1. Click on the **Virtual Drives** link on the WebBIOS main menu.

The Virtual Drives dialog appears, as shown in [Figure 145](#).

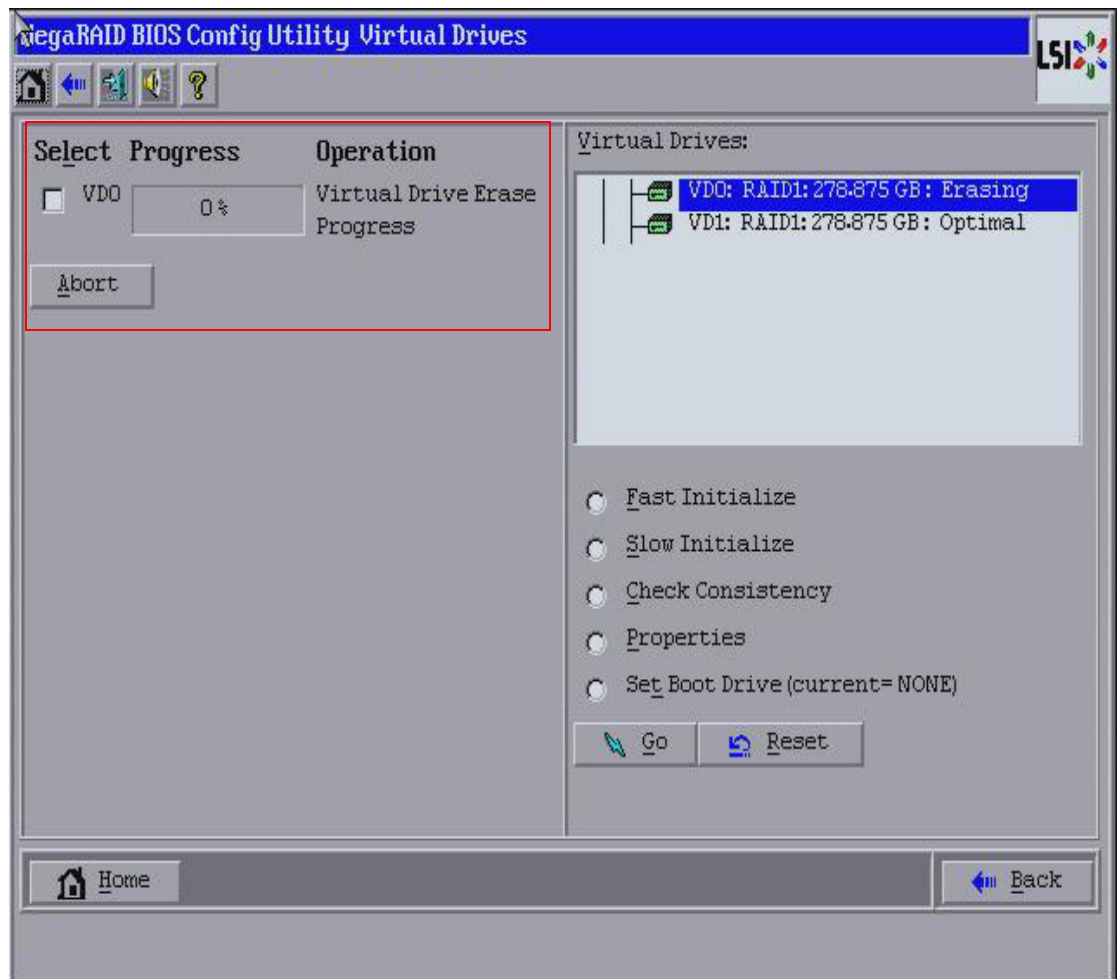


Figure 145: Virtual Drive Dialog

2. To abort the virtual drive erase, select the check box of the operation you want to abort, click **Abort**.

4.13 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 configuration utility to view these event messages. To do this, click **Events** on the main WebBIOS configuration utility dialog. The Event Information dialog appears, as shown in [Figure 146](#).

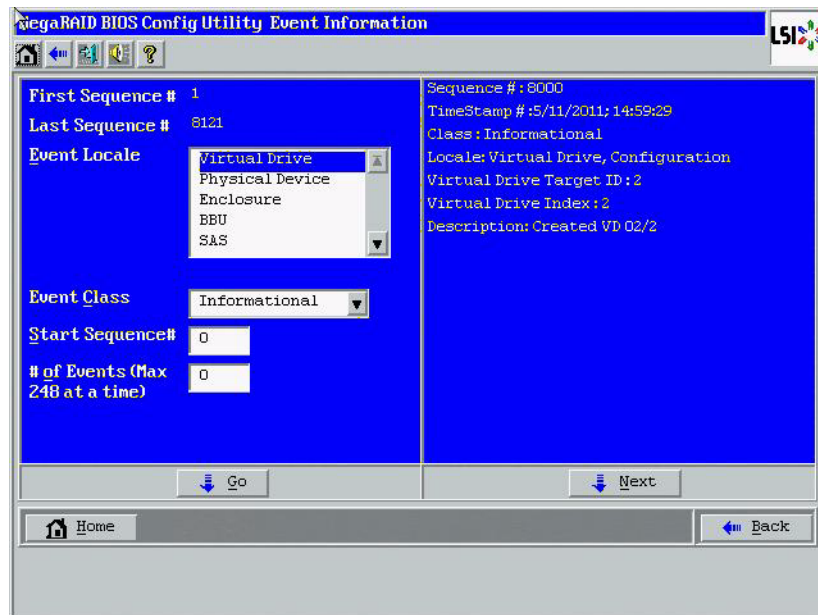


Figure 146: Event Information Dialog

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

To view event information, follow these steps:

1. Select an event locale from the **Event Locale** drop-down list. 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 the 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** to page forward or **Prev** to page backward through the sequence of events.
6. Optionally, select different event criteria in the left panel, and click **Go** again to view a different sequence of events.

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

4.14 Managing Configurations

4.14.1 Running a Consistency Check

This section includes information about maintaining and managing storage configurations.

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 configuration utility main dialog, select a virtual drive.
2. Click **Virtual Drives**.
3. When the Virtual Drive dialog appears, select **CC** in the lower-left panel, and click **Go**.

The consistency check begins.

If the WebBIOS configuration utility 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 data might be corrupted.

4.14.2 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 configuration utility 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.

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

To delete a virtual drive, follow these steps.

1. Access the Virtual Drive dialog by clicking on a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.

The Virtual Drive dialog appears.

2. Select **Delete** in the bottom panel under the heading Operations, and click **Go**.
3. When the message appears, confirm that you want to delete the virtual drive.

If a virtual drive is associated with a CacheCade virtual drive with a write policy, the following confirmation screen appears.

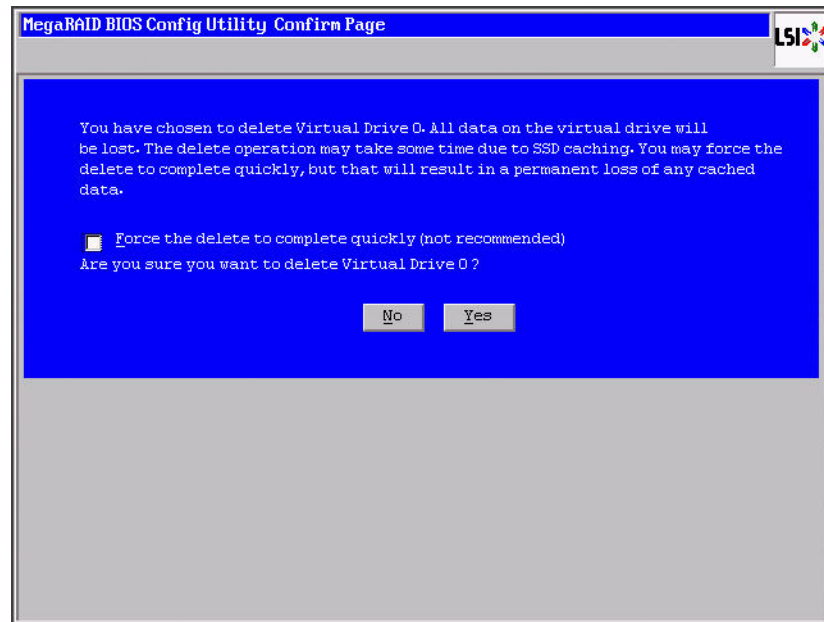


Figure 147: WebBIOS CU Confirmation Screen

Click **Yes** to delete the virtual drive.

NOTE: You may select the **Force the delete to complete quickly** check box to quickly complete the delete operation. It is however, not recommended to perform this action.

4.14.3 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 WebBIOS 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, do *not* appear. To use drives with existing configurations, you must first clear the configuration on those drives.

If WebBIOS configuration utility detects a foreign configuration, the Foreign Configuration dialog appears, as shown in [Figure 148](#).

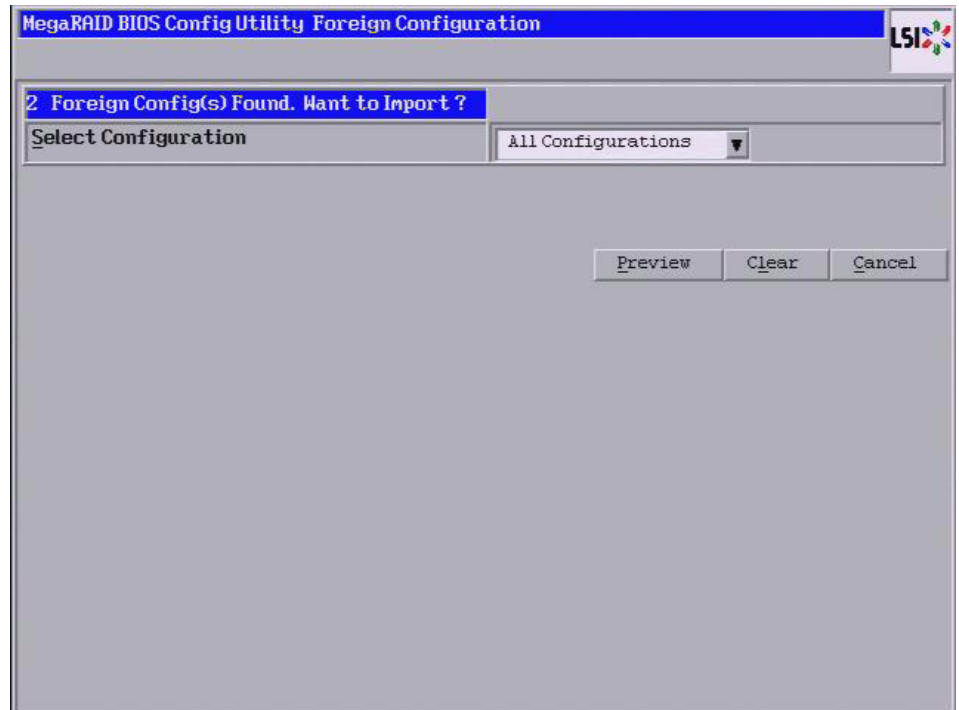


Figure 148: Foreign Configuration Dialog

Follow these steps to import or clear a foreign configuration.

1. Click the drop-down list to show the configurations.
The GUID (Global Unique Identifier) entries on the drop-down list are OEM names and will vary from one installation to another.
2. Either select a configuration, or select **All Configurations**.
3. Perform one of the following steps:
 - Click **Preview** to preview the foreign configurations.
The Foreign Configuration Preview dialog appears, as shown in [Figure 149](#).
 - Click **Clear** to clear the foreign configurations and reuse the drives for another virtual drive.

If you click **Cancel**, it cancels the importation or preview of the foreign configuration.

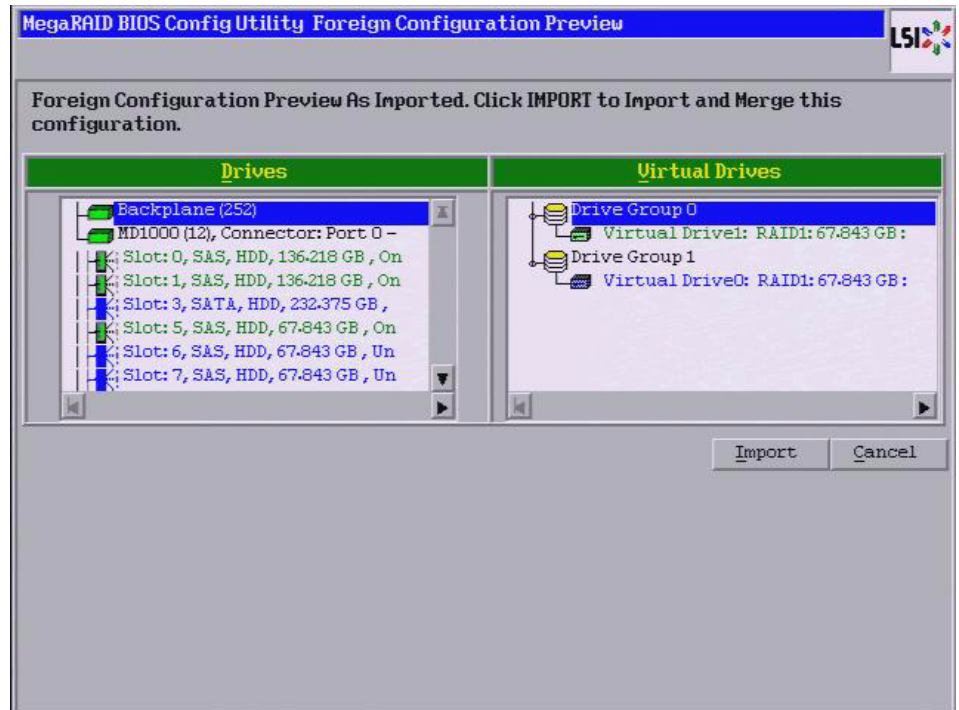


Figure 149: Foreign Configuration Preview Dialog

The right panel shows the virtual drive properties of the foreign configuration. In this example, there are two RAID 1 virtual drives with 67.843 GB each. The left panel shows the drives in the foreign configuration.

4. Click **Import** to import these foreign configurations and use them on this controller. If you click **Cancel**, you return to the dialog shown in [Figure 148](#).

4.14.3.1 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 dialog to import or clear the foreign configuration in each case. The import procedure and clear procedure are described in [Section 4.14.3, Importing or Clearing a Foreign Configuration](#).

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

NOTE: 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.

- 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 4.14.1, Running a Consistency Check](#) for more information about checking data consistency.

- 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 4.14.1, Running a Consistency Check](#) for more information about checking data consistency.

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

4.14.3.2 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 website at: <http://www.lsi.com/cm/DownloadSearch.do>.

4.14.3.3 Troubleshooting Information

An IR virtual drive can have either 64 MB or 512 MB available for metadata at the end of the drive. This data is in LSI Data Format (LDF). MegaRAID virtual drives have 512 MB 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 MB in the metadata, which is the same amount of mega data as in a MegaRAID virtual drive. If the IR virtual drive has only 64 MB when you attempt to import it into MegaRAID, the import will fail because the last 448 MB of your data will be overwritten and the data lost.

If your IR virtual drive has only 64 MB 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.

To import an IR virtual drive into a MegaRAID system, use the Foreign Configuration Preview dialog to import or clear the foreign configuration. The import procedure and the clear procedure are described in [Section 4.14.3, *Importing or Clearing a Foreign Configuration*](#).

4.14.4 Importing Foreign Configurations

After you create a security key, you can run a scan for a foreign configuration and import a locked configuration. (You can import unsecured or unlocked configurations when security is disabled.) A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the WebBIOS configuration utility to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See [Section 4.14.3, *Importing or Clearing a Foreign Configuration*](#) for the procedures used to import or clear a foreign configuration.

To import a foreign configuration, you must first enable security to allow importation of locked foreign drives. If the drives are locked and the controller security is disabled, you cannot import the foreign drives. Only unlocked drives can be imported when security is disabled.

After you enable the security, you can import the locked drives. To import the locked drives, you must provide the security key used to secure them. Verify whether any drives are left to import as the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all of the drives are imported, there is no configuration to import.

4.14.5 Import Foreign Drives in EKM/EKM Secured Locked Drives

When you scan a foreign configuration, and if any of the foreign EKM secured locked drives are present when EKM is enabled, the following wizard is displayed.



Figure 150: Import Foreign Drives in EKM Mode

4.14.6 Importing Foreign Drives for LKM-Secured Locked Drives

When you scan a foreign configuration, and if any of the foreign LKM-secured, locked drives prompt you to enter a security key, assuming that the drive security is enabled on the controller, the following wizard appears.

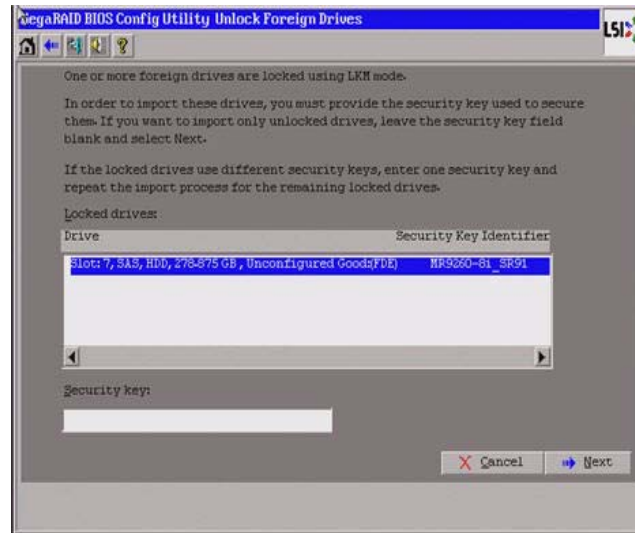


Figure 151: Import Foreign Drives in EKM Mode

- The **Locked drives** field displays the locked physical drives and their corresponding security key identifiers.
- You must enter the drive security key in the **Security Key** field.

4.14.7 Importing Foreign Drives in LKM Mode EKM-Secured Locked Drives

When you scan a foreign configuration, and if you import foreign drives in LKM mode, when EKM-secured locked drives are present, the following warning appears.



Figure 152: Import Foreign Drives in LKM Mode

4.14.8 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 restart 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 configuration utility 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 23 lists the number of additional drives required when you change the RAID level of a virtual drive.

Table 23: 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 6	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 configuration utility main dialog, select **Virtual Drives**.
2. Choose your virtual drive from the list. If only one virtual drive is configured, you will automatically be taken to the **Virtual Drives** menu.

3. From the Virtual Drives menu, select **Properties**.
4. From the **Properties** menu, select **Adv Opers** under the **Advanced Operations** heading.

The Advanced Operations dialog appears, as shown in [Figure 153](#).

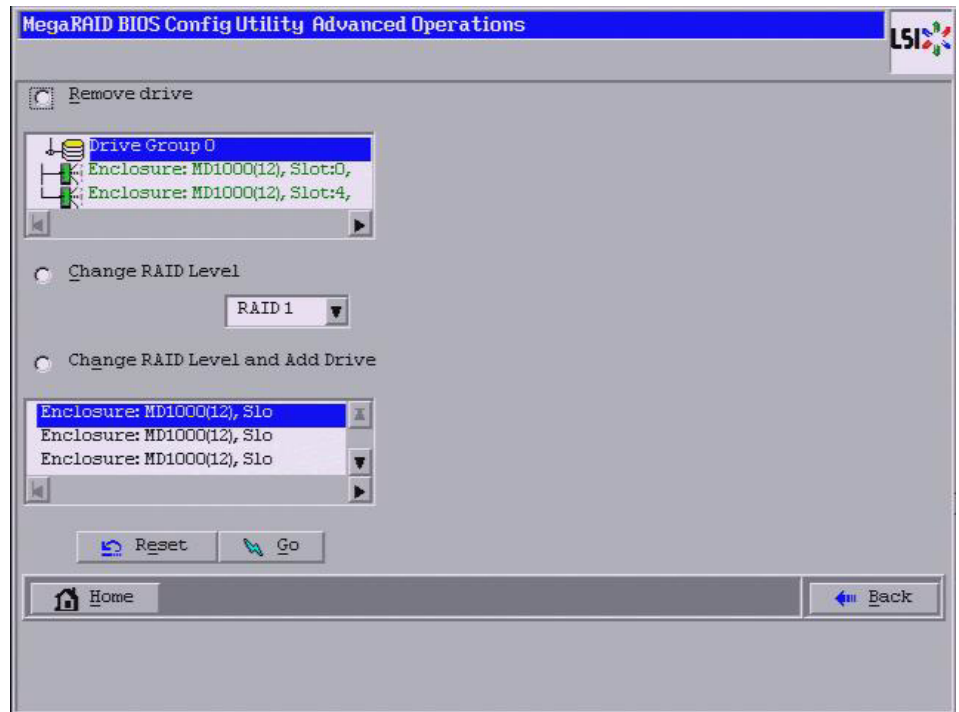


Figure 153: Advanced Operations Dialog

5. Select either **Change RAID Level** or **Change RAID Level and Add Drive**.
 - If you select **Change RAID Level**, change the RAID level from the drop-down list.
 - If you select **Change RAID Level and Add Drive**, change the RAID level from the drop-down list, and select one or more drives to add from the list of drives.

The available RAID levels are limited, based on the current RAID level of the virtual drive plus the number of drives available.

6. Click **Go**.
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 configuration utility.

4.14.9 New Drives Attached to a MegaRAID Controller

When you insert a new drive on a MegaRAID system, if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD. If the drive does contain valid DDF metadata, its drive state is Unconfigured Good.

A new drive in JBOD drive state is exposed to the host operating system as a stand-alone drive. Drives in JBOD drive state are not part of the RAID configuration because they do not have valid DDF records. The operating system can install and run anything on JBOD drives.

Automatic rebuilds always occur when the drive slot status changes, for example, when you insert a drive or remove a drive, so that a hot spare can be used. However, a new drive in JBOD drive state (without a valid DDF record), will not perform an automatic rebuild.

To start an automatic rebuild on the new JBOD drive, you must change the drive state from JBOD to Unconfigured Good. (Rebuilds start only on Unconfigured Good drives.) After you set the drive state to Unconfigured Good, the drive state information always remains on the drive, and you can use the drive for configuration.

See [Section 4.14.3.3, *Troubleshooting Information*](#) for more information about DDF and metadata.

4.15 WebBIOS Dimmer Switch

This section describes changing the power-save settings using the Dimmer Switch feature.

The power savings is done by reducing power consumption of drives that are not in use by spinning down the unconfigured drives, hot spares, and configured drives.

Perform the following steps to change the Dimmer switch feature.

1. Select the **Controller Properties** option from the WebBIOS main dialog.

The Controller Information dialog appears, as shown in the following figure.



Figure 154: Controller Information Dialog - Power Save Mode for Config Drives

2. Click **Next**.

- Click **Next** again. You will go to the third Controller properties page.
The Controller Properties dialog appears as shown in the [Figure 155](#).

MegaRAID BIOS Config Utility Controller Properties			
Properties			
Battery Backup	Present	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Disabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	4
Reconstruction Rate	30	Spinup Delay	2
NCQ	Enabled		
<input type="button" value="Submit"/> <input type="button" value="Reset"/> <input type="button" value="Next"/>			
<input type="button" value="Home"/>		<input type="button" value="Back"/>	

Figure 155: Controller Properties Dialog

- Click **Next**.

The Controller Properties dialog with the **Manage Powersave** option appears as shown in the following figure.

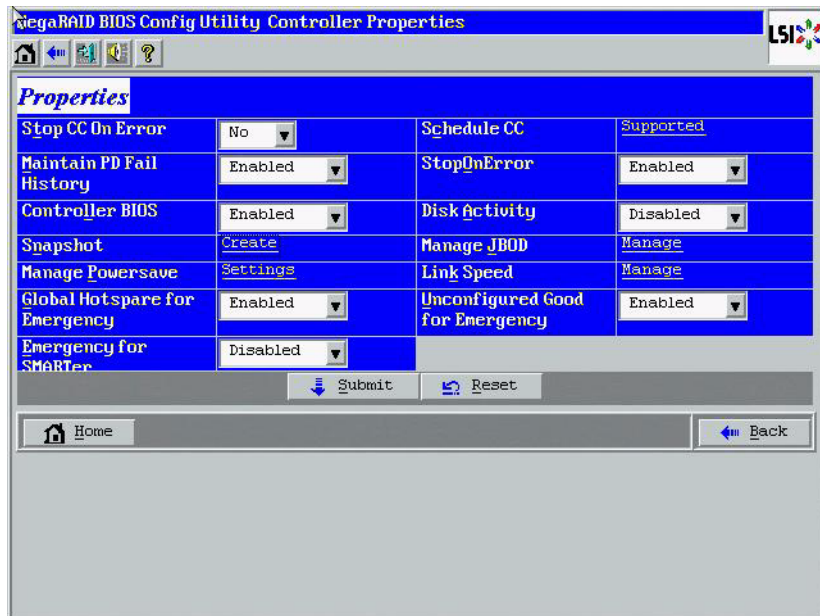


Figure 156: Controller Properties - Manage Powersave Option Dialog

5. Click the **Manage Powersave** option.

The **Power Save Setting -- Specify Power Save Setting** dialog appears, as shown in the following figure.

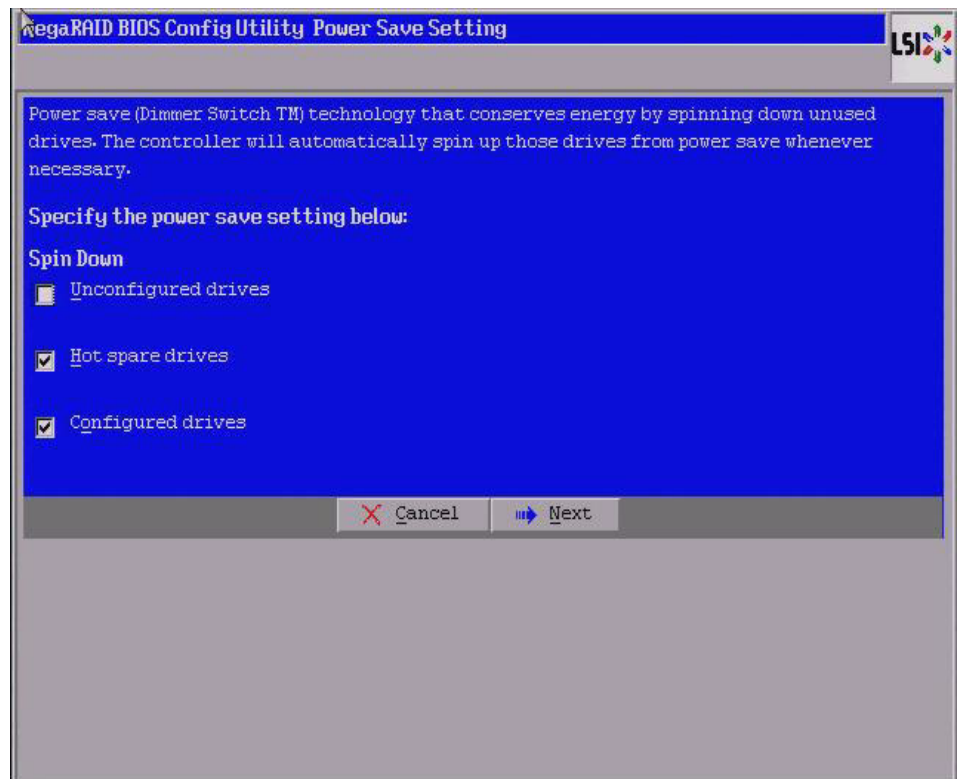


Figure 157: Power Save Setting Dialog - Specify Power Save Setting

6. Select the following check boxes:
 - Select the **Unconfigured drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
 - Select the **Hot spare drives** check box to let the controller enable the hot spare drives to enter the Power-Save mode.
 - Select the **Configured drives** check box to let the controller enable the Configured drives to enter the Power-Save mode.

7. Click **Next**. The **Power Save Setting-- Power Save Mode** dialog appears.

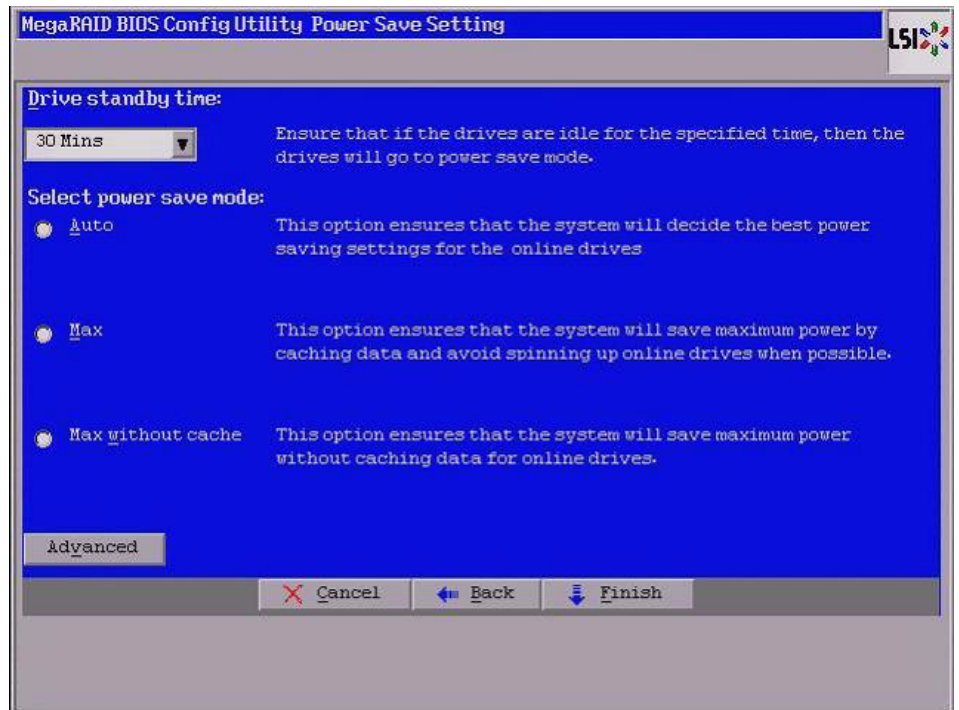


Figure 158: Power Save Setting - Power Save Mode

8. Click **Finish**.

The Confirmation message appears, as shown in [Figure 160](#).

In the **Power-Save Setting -- Specify Power Save Setting** dialog ([Figure 157](#)), if you select all the check boxes, except the configured drives, then the following dialog appears.



Figure 159: Power Save Setting - Except Configured Drive Dialog



Figure 160: Confirm Page Dialog - Confirmation Message

If you do not select any option in the **Power Save settings --- Specify Power Save Setting** dialog (Figure 157) and click **Next**, the following message appears.

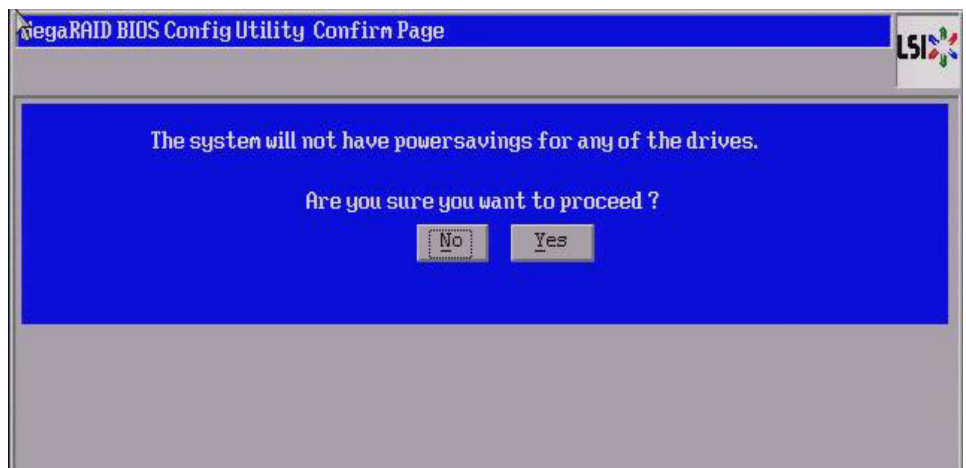


Figure 161: Power-Save Settings Not Saved Message

4.15.1 Power-Save Mode

You can select the drive standby time and the Power-Save mode by selecting the **Auto**, **Max**, and **Max without cache** options in Figure 158.

1. Select the drive standby time using the drop-down list.
2. Select the power save mode by selecting one of the radio buttons.
3. Click **Finish**.

4.15.2 Power Save Settings-Advanced

You can schedule the drive active time by selecting the start time and end time in the Power-Save Setting dialog.

Perform the following steps to schedule the drive active time.

1. Click the **Advanced** button in the Power-Save setting dialog as shown in [Figure 158](#).

The Power-Save Settings Advanced dialog appears, as shown in the following figure.

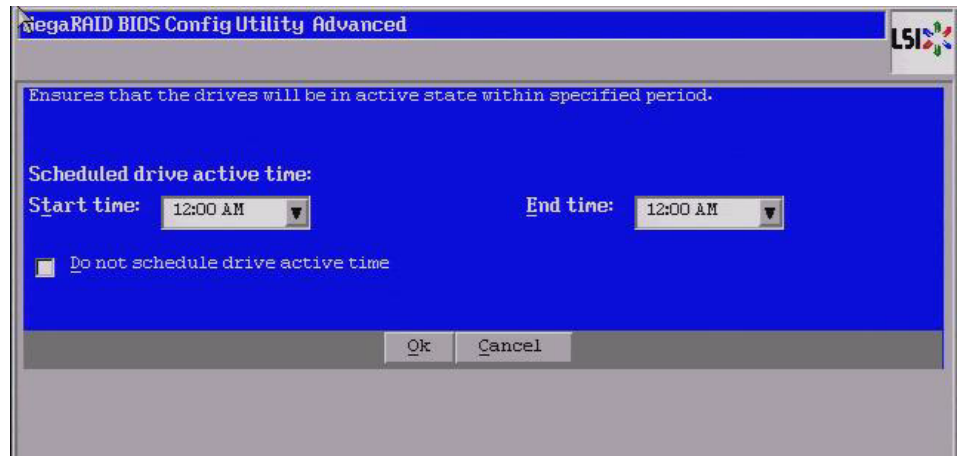


Figure 162: Power - Save Settings Advanced

2. Select the start time and end time from the **Scheduled drive active time** field using the **Start time** and **End time** drop-down list.
3. Click **OK**.

The drive active time is scheduled.

NOTE: Select the **Do not schedule drive active time** check box if you do not want to schedule the drive active time.

4.15.3 Power-Save While Creating Virtual Drives

You can select the power saving policy while creating virtual drives using the **Power save Mode** drop-down list.

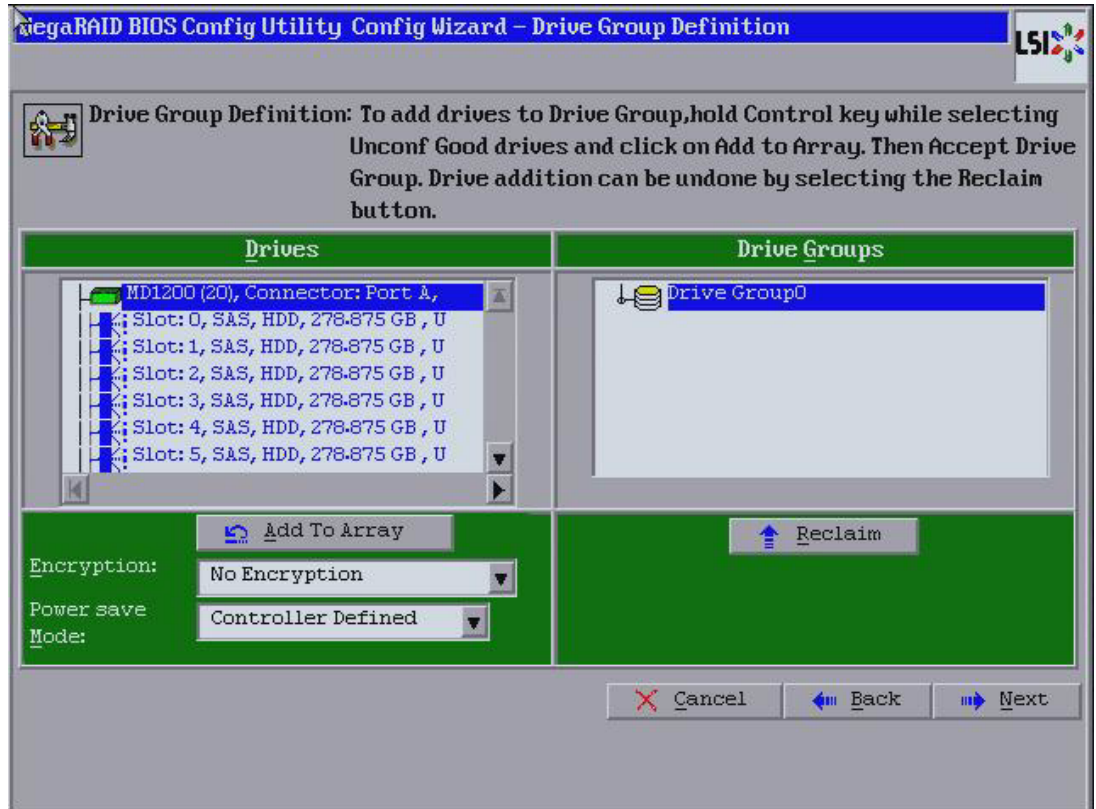


Figure 163: Power Save Mode while Creating Virtual Drives

The power save mode can be **Max**, **Max without cache**, **Auto**, **None**, and **Controller defined**.

Chapter 5

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 the MegaRAID SAS RAID controllers and the devices connected to them.

NOTE: The CT supports only the MegaRAID controllers that support SAS and SATA II. It does not support other types of MegaRAID controllers, such as U320, SATA I, or IDE.

NOTE: The IA-64 release for the Windows® operating system is similar to the 32-bit release, so you can follow the 32-bit instructions. Applications that are 32-bit that were validated on an x64 system, such as the Intel Market system, can use the 32-bit instructions, also.

5.1 Product Overview

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

- Configure the 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
- Change power setting (dimmer switch)
- 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 MS-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 the firmware event logs
- Display help for how to use the command line options
- Enable or disable snapshots (for the Recovery advanced software feature)
- Create and delete snapshots and views of a virtual drive
- Roll back the virtual drive to an older snapshot
- Display snapshot properties
- Create a CacheCade 2.0 SSD Read Caching virtual drive to use as secondary cache
- Display battery CacheCade 2.0 SSD Read Caching 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 error messages are listed in [Appendix B](#).

5.2 Novell NetWare, SCO, Solaris, FreeBSD, and MS-DOS Operating System Support

The MegaCLI Configuration Utility functions under the Novell® NetWare®, SCO® OpenServer™, SCO UnixWare®, Solaris®, FreeBSD®, and MS-DOS operating systems in the same way that it does under the Windows and Linux® operating systems. All 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 disk drive. The image file name is `MegaCLI.image`. The disk is provided so that you can distribute MegaCLI and install the executable file later as needed.

For the Solaris operating system, LSI releases MegaCLI as a package that can be installed like any other package installation in Solaris.

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.

5.3 Command Line Abbreviations and Conventions

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

5.3.1 Abbreviations Used in the Command Line

The following table lists the abbreviations for the virtual drive parameters used in the following sections.

Table 24: Command Line Abbreviations

Abbreviation	Description
WB	Write Back write policy
WT	Write Through 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

5.3.2 Conventions

You can specify multiple values for some options. You can enter commands for a single controller (`-aN`), multiple controllers (`-a0, 1, 2`) or work on all present controllers (`-aALL`). The options are denoted as `-aN | -a0, 1, 2 | -aALL` in this document and specify 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.

The following table describes the conventions used in the options.

Table 25: 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 25: Conventions (Continued)

Convention	Description
[E0 : S0 , E1 , S1 , ...]	Specifies when one or more physical devices need to be specified in the command line. Each [E : S] pair specifies one physical device, where E means the device ID of the enclosure in which a drive resides, and S means the slot number of the enclosure. 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.
[]	Indicates that the parameter is optional except when it is used to specify physical devices. For example, [WT] means the write policy (Write Through) is optional. If you enter WT at the command line, the application will use Write Through write policy for the virtual drive. Otherwise, it uses the default value for the parameter.
{ }	Indicates that the parameters are grouped and that they must be given at the same time.
-Force	Specifies that the MegaCLI utility does not ask you for confirmation before it performs this command. You might lose data using this option with some commands.

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 dialog.

5.4 Pre-boot MegaCLI

A second CLI utility, known as Pre-boot MegaCLI (PCLI), is available. You can enter this utility during bootup. PCLI gives you an alternative way to access the MegaCLI utility.

To access PCLI, while the host computer is booting, hold down the **Ctrl** key and press the **Y** key when the following text appears on the dialog:

Copyright© LSI Logic Corporation

Press <Ctrl><Y> for Preboot CLI

The following commands that are in the regular MegaCLI utility are not available in PCLI:

- AdpSetVerify
- AdpCcSched
- AdpDiag
- AdpBatTest
- option ProgDsply
- CfgSave
- CfgRestore
- AdpBbuCmd

- AdpFacDefSet
- AdpFwFlash
- AdpGetConnectorMode
- AdpSetConnectorMode
- DirectPdMapping
- ShowEnclList
- ShowVpd
- EnclLocate
- PdFwDownload
- SetFacDefault
- PDCpyBk
- AdpFwDump
- Snapshot Enbl/Setprop/Dsbl/TakeSnapshot/DeleteSnapshot
CreateView/DeleteView/Info/Clean/GetViewInfo
- AdpSetProp DefaultSnapshotSpace
DefaultViewSpace/AutoSnapshotSpace

5.5 CacheCade Related Options

Use the commands in this section to perform actions related to CacheCade 2.0 SSD Read Caching software and CacheCade Pro 2.0 SSD Read/Write Caching software.

CacheCade 2.0 SSD Read Caching improves application performance by expanding the MegaRAID read-caching capacity and introducing write caching. The CacheCade 2.0 SSD Read Caching feature uses high-performing CacheCade 2.0 SSD Read Caching software as a secondary tier of cache to provide faster reads and to maximize transactional I/O performance.

Using CacheCade 2.0 SSD Read Caching software as controller cache allows for very large data sets to be present in cache, delivering performance up to 50 times greater than regular cache in read-intensive applications, such as online transaction processing (OLTP) and file and Web server workloads. The solution is designed to accelerate the I/O performance of HDD-based drive groups while only requiring a small investment in CacheCade 2.0 SSD Read Caching software technology.

To support full-throughput for multiple direct-attached CacheCade 2.0 SSD Read Caching software, this feature reduces I/O-processing overhead in the SAS 6Gb/s MegaRAID controllers. CacheCade 2.0 SSD Read Caching offers performance equivalent to flash-based controllers and better performance for RAID 5 and RAID 6 when compared to Fusion I/O.

5.5.1 Create a Solid State Drive Cache Drive to Use as Secondary Cache

Use the command in the following table to create a cache drive using CacheCade 2.0 SSD Read Caching software. You can use that cache as secondary cache. CacheCade 2.0 SSD Read Caching software has much greater capacity than HDDs.

Table 26: Create a Solid State Cache Drive to Use as Secondary Cache

Convention	MegaCli -CfgCacheCadeAdd -Physdrv[E0:S0,...] {-Name LdNamestring} -aN -a0,1,2 -aALL
Description	This command is used to create CacheCade software that you can use as secondary cache. -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots to use to construct a drive group. -Name LdNamestring: This is the name given to the CacheCade software cache drive.

5.5.2 Delete a Solid State Drive Cache Drive

Use the command in the following table to delete a CacheCade 2.0 SSD Read Caching software cache drive or multiple cache drives on the selected controllers.

Table 27: Delete Solid State Cache Drives

Convention	MegaCli -CfgCacheCadeDel -Lx -L0,2,5... -LALL -aN -a0,1,2 -aALL
Description	Deletes the specified CacheCade software cache drive or drives on the selected controllers. You can delete multiple CacheCade software cache drives or all of the CacheCade software caches.

5.5.3 Associate/Disassociate Virtual Drives

Use this command in the following table to associate or disassociate virtual drives with a CacheCade Pro 2.0 virtual drive.

Table 28: Associate/Disassociate Virtual Drives

Convention	MegaCLI -Cachecade -assign -remove -Lx -L0,1,2 -LALL -aN -a0,1,2 -aALL
Description	Assigns or removes association of virtual drives with the CacheCade pool. -assign: associates virtual drives with a CacheCade Pro 2.0 virtual drive. -remove: disassociates virtual drives with a CacheCade Pro 2.0 virtual drive.

5.5.4 Display CacheCade Pro 2.0 Configurations on a Controller

Use this command in the following table to display all the existing CacheCade Pro 2.0 configurations on a selected controller.

Table 29: Display CacheCade Pro 2.0 Configurations on a Controller

Convention	MegaCLI -CfgCacheCadeDsply -aN -a0,1,2 -aALL
Description	Displays all the existing CacheCade Pro 2.0 configurations on a selected controller.

5.5.5 Create a RAID Drive Group for CacheCade Pro 2.0 from All Unconfigured Good Drives

Use the command in the following table to create one RAID drive group, for CacheCade Pro 2.0, 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: The firmware supports only 32 drives per drive group. If there are more than 32 unconfigured good drives, MegaCLI cannot configure any of the drives, and the command will abort.

Table 30: Create a RAID Drive Group for CacheCade Pro 2.0 from All Unconfigured Good Drives

Convention	<pre>MegaCLI -CfgLdAdd -rX[E0:S0,E1:S1,...] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-szXXX [-szYYY ...]] [-strpszM] [-Hsp[E0:S0,...]] [-AfterLdX] [-Force] [FDE CtrlBased] [-Default -Automatic -None -Maximum -MaximumWithoutCaching] [-Cache] -aN</pre>
Description	<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, 1, 5, or 6. 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>Hsp: Specifies drive to make the hot spare with.</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> <p>-Cache : Specifies that SSD Caching is enabled for the Drive group.</p> <p>NOTE: Previously, -szXXX expressed capacity in MB but now you can enter the capacity in your choice of units. For example, to create a virtual drive of 10 GB, enter the size as sz10GB. If you do not enter a unit, by default it is considered as MB.</p>

5.5.6 Remove Blocked Access on a Virtual Drive

Use this command in the following table to change the access policy for a virtual drive by removing a blocked access on that virtual drive. At times, an error may occur in the CacheCade Pro 2.0 virtual drive and this causes a blocked access to the associated virtual drive.

Table 31: Remove Blocked Access on a Virtual Drive

Convention	MegaCLI -LDSetProp {-Name LdNamestring} -RW RO Blocked RemoveBlocked WT WB ForcedWB [-Immediate] RA NORA ADRA -aN -a0,1,2 -aALL
Description	Allows you to change the following virtual drive parameters: [WT WB ForcedWB] : Specifies the write policy- WT (Write Through, WB (Write Back), ForcedWB (Forced Write Back). If you specify the ForcedWB parameter, the write policy will always be writeback, even if the virtual drive becomes degraded. -Immediate: Indicates that the changes take place immediately. -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. -RemoveBlocked: Removes the blocked access on the associated virtual drive. -EnDskCache: Enables drive cache. -DisDskCache: Disables drive cache.

5.5.7 Create RAID 0 Configuration with SSD Caching

Use this command in the following table to create virtual drives with RAID 0 configurations and enable SSD caching on them

Table 32: Create RAID 0 Configuration with SSD Caching

Convention	MegaCLI -CfgEachDskRaid0 [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-strpszM] [FDE CtrlBased] [-Default -Automatic -None -Maximum -MaximumWithoutCaching] [-Cache] -aN -a0,1,2 -aALL
Description	Creates virtual drives with RAID 0 and enables SSD caching on these newly created virtual drives. -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. [-Default -Automatic -None -Maximum -MaximumWithoutCaching]: If the controller supports power savings on virtual disk, these options specify the possible levels of power savings that can be applied on a virtual disk. [-Cache] : Specifies that SSD caching is enabled.

5.5.8 Create a RAID Level 10, 50, 60 (spanned) Configuration with SSD Caching

Use the command in the following table to create a RAID 10, RAID 50, or RAID 60 configuration with SSD caching to the existing configuration on the selected controller.

Table 33: Create a RAID Level 10, 50, 60 (spanned) Configuration with SSD Caching

<p>Convention</p>	<pre>MegaCLI -CfgSpanAdd -r10 -Array0[E0:S0,E1:S1] -Array1[E0:S0,E1:S1] [-ArrayX[E0:S0,E1:S1] ...] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-szXXX[-szYYY ...]] [-strpszM] [-AfterLdX] [-Force] [FDE CtrlBase d] [-Default -Automatic -None -Maximum -MaximumWithoutCaching]] [-Cache] -aN</pre>
<p>Description</p>	<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 X 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>-szXXXXXXXXXX</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 <code>afterLd</code> option is specified but the controller does not support slicing in the spanned drive groups.</p> <hr/> <p>NOTE: Previously, <code>-szXXX</code> expressed capacity in MB but now you can enter the capacity in your choice of units. For example, to create a virtual drive of 10 GB, enter the size as <code>sz10GB</code>. If you do not enter a unit, by default it is considered as MB.</p> <hr/> <p><code>[-Default -Automatic -None -Maximum -MaximumWithoutCaching]</code> : If the controller supports power savings on virtual disk, these options specify the possible levels of power savings that can be applied on a virtual disk.</p> <p><code>[-Cache]</code> : Specifies that SSD caching is enabled.</p>

5.5.9 Delete Virtual Drives with SSD Caching

Use the command in the following table to delete one or all virtual drives that have SSD caching enabled.

Table 34: Delete Virtual Drives with SSD Caching

Convention	MegaCLI -CfgLdDel -LX -L0,2,5... -LALL [-Force] -aN -a0,1,2 -aALL
Description	Deletes virtual drives associated with a CacheCade virtual drive, with a write policy, using the Force option. [-Force] Specifies that the data is not flushed before deleting the virtual drive.

5.5.10 Clear Configurations on CacheCade Pro 2.0 Virtual Drives

Use the command in the following table if the selected controller has any CacheCade Pro 2.0 virtual drives or if any data exists in the cache and you want to clear all existing configurations on a controller.

Table 35: Clear Configurations on CacheCade Pro 2.0 Virtual Drives

Convention	MegaCLI -cfgclr [-force] -a0
Description	Clears all existing configurations on a controller if the controller has any CacheCade 2.0 virtual drives or if any data exists in the cache. [-Force] Specifies that the data is not flushed before deleting the virtual drive.

5.5.11 Create a CacheCade Pro 2.0 Virtual Drive with RAID level and Write Policy

Use the command in the following table to create a CacheCade Pro 2.0 virtual drive with RAID level and Write Policy settings.

Table 36: Create a CacheCade Pro 2.0 Virtual Drive with RAID level and Write Policy

Convention	MegaCLI -CfgCacheCadeAdd [-rX] -Physdrv[E0:S0,...] {-Name LdNamestring} [WT WB ForcedWB] [-assign -LX L0,2,5... LALL] -aN -a0,1,2 -aALL
Description	-Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots to use to construct a drive group. -Name LdNamestring: This is the name given to the CacheCade software cache drive. [-rX]: Specifies the RAID level. [WT WB ForcedWB]: Specifies the write policy- WT (Write Through, WB (Write Back), ForcedWB (Forced Write Back). While creating a CacheCade virtual drive, if you specify the ForcedWB parameter, even if the CacheCade virtual drive goes into a degraded mode, the write policy stays as Write Back. [-assign]: Specifies that the virtual drive can be associated with a CacheCade virtual drive.

5.6 Software License Key

Use the commands in this section to obtain a software license key to enable the advanced features present in the controller.

Table 37: Software License Key

Convention	MegaCli ELF -GetSafeId -a0
Description	Displays the Safe ID of the controller.
Convention	MegaCli ELF -ControllerFeatures -aN -a0,1,2 -aALL

Table 37: Software License Key

Description	Displays the Advanced Software Options that are enabled on the controller including the ones in trial mode.
Convention	MegaCli -ELF -Applykey key <-val> [Preview] -aN -a0,1,2 -aALL
Description	Applies the Activation Key either in preview mode or in real mode.
Convention	MegaCli -ELF -TransferToVault -aN -a0,1,2 -aALL
Description	Transfers the Activated Advanced Software Options from NVRAM to keyvault.
Convention	MegaCli -ELF -DeactivateTrialKey -aN -a0,1,2 -aALL
Description	Deactivates the trial key.
Convention	MegaCli -ELF -ReHostInfo -aN -a0,1,2 -aALL
Description	Displays the re-host information, and if re-hosting is necessary it displays the controller and keyvault serial numbers.
Convention	MegaCli -ELF -ReHostComplete -aN -a0,1,2 -aALL
Description	Indicates to the controller that re-host is complete.

5.7 SafeStore Security Options

Use the commands in this section to manage the SafeStore Security feature. This feature offers the ability to encrypt data on disks and use disk-based key management to provide data security. With this feature, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive level.

This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting disks, if you remove a drive from its storage system or the server in which it is housed, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

Any encryption solution requires management of the encryption keys. This feature provides a way to manage these keys. You can change the encryption key for all ServeRAID controllers that are connected to SED drives. All SED drives, whether locked or unlocked, always have an encryption key. This key is set by the drive and is always active. When the drive is unlocked, the data to host from the drive (on reads) and from the host to the drive cache (on writes) is always provided. However, when resting on the drive platters, the data is always encrypted by the drive.

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

See [Chapter 3, SafeStore Disk Encryption](#) for more information about the SED feature.

5.7.1 Use Instant Secure Erase on a Physical Drive

Use the command in the following table to perform an instant secure erase of data on a physical drive. The Instant Secure Erase feature lets you erase data on SED drives.

Table 38: Use Instant Secure Erase on a Physical Drive

Convention	MegaCli -PDInstantSecureErase -PhysDrv[E0:S0,E1:S1,...] [-Force] -aN -a0,1,2 -aALL
Description	Erases the data on a specified drive or drives. -PDInstantSecureErase: Use the Instant Secure Erase feature to erase data on a drive or drives. -PhysDrv[E0:S0,...]: Specifies the drives on which you want to perform the Instant Secure Erase. -Force: Specifies that the MegaCLI utility does not ask you for confirmation before it performs this command (you might lose data using this option with some commands).
	NOTE: NOTE: Previously -szXXX expressed capacity in MB, but now you can enter the capacity in your choice of units. For example, to create a virtual drive of 10 GB, enter the size as sz10GB. If you do not enter a unit, MB is used as the default unit.

5.7.2 Secure Data on a Virtual Drive

Use the command in the following table to secure data on a virtual drive.

Table 39: Secure Data on a Virtual Drive

Convention	MegaCli -LDMakeSecure -Lx -L0,1,2,... -Lall -aN -a0,1,2 -aALL
Description	Secures data on a specified virtual drive or drives.

5.7.3 Destroy the Security Key

Use the command in the following table to destroy the security key.

Table 40: Destroy the Security Key

Convention	MegaCli -DestroySecurityKey [-Force] -aN
Description	Destroys the security key. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For SED drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted.

5.7.4 Create a Security Key

Use the command in the following table to create a security key.

Table 41: Create a Security Key

Convention	MegaCli -CreateSecurityKey -SecurityKey ssssssssss [-Passphrase ssssssssss] [-KeyID kkkkkkkkkk] -aN
Description	<p>Creates a security key based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. If the security key is unavailable, user data is irretrievably lost. You must take all precautions to never lose the security key.</p> <p>-CreateSecurityKey: Creates the security key.</p> <p>-SecurityKey ssssssssss: Enters the new security key. The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.</p> <p>[-Passphrase ssssssssss]: Enters the new passphrase. The passphrase is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.</p>

5.7.5 Create a Drive Security Key

If you want to use the security key using electronic key management system (EKMS), the EKMS must provide the security key. You can create a security key using EKMS, or switch from EKM to LKM, or from LKM to EKM.

Table 42: Drive Security Key

Convention	MegaCli -CreateSecurityKey useEKMS -aN
Description	Creates security key using EKMS.
Convention	MegaCli -ChangeSecurityKey -SecurityKey ssssssssss [-Passphrase ssssssssss] [-KeyID kkkkkkkkkk] -aN
Description	To change the security from EKMS to LKM.
Convention	MegaCli -ChangeSecurityKey useEKMS -OldSecurityKey ssssssssss -aN
Description	To change security from LKM to EKM.
Convention	MegaCli -ChangeSecurityKey -useEKMS -aN-
Description	Rekeying in EKMS

5.7.6 Change the Security Key

Use the command in the following table to change they security key to a new security key.

Table 43: Change the Security Key

Convention	<code>MegaCli -ChangeSecurityKey -OldSecurityKey sssssssssss -SecurityKey sssssssssss [-Passphrase sssssssssss] [-KeyID kkkkkkkkkkk] -aN</code>
Description	<p>Changes a security key to a new security key.</p> <p><code>-ChangeSecurityKey</code>: Changes the security key.</p> <p><code>-OldSecurityKey sssssssssss</code>: Enters the old security key. The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.</p> <p><code>-SecurityKey sssssssssss</code>: Enters the new security key. The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.</p> <p><code>[-Passphrase sssssssssss]</code>: Enters the new passphrase. The pass phrase is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.</p> <p><code>[-KeyID kkkkkkkkkkk]</code>: Enters the security key ID. The key ID displays when you have to enter a security key. If you have multiple security keys, the security key ID helps you determine which security key to enter.</p>

5.7.7 Get the Security Key ID

Use the command in the following table to display the security key ID.

Table 44: Get the Security Key ID

Convention	<code>MegaCli -GetKeyID [-PhysDrv[E0:S0]] -aN</code>
Description	<code>-GetKeyID</code> : Displays the security key ID.

5.7.8 Set the Security Key ID

Use the command in the following table to set the security key ID.

Table 45: Set the Security Key ID

Convention	<code>MegaCli -SetKeyID -KeyID kkkkkkkkkkk -aN</code>
Description	<p><code>-SetKeyID</code>: Set the security key ID.</p> <p><code>-KeyID kkkkkkkkkkk</code>: Enters the security key ID. The key ID displays when you have to enter a security key. If you have multiple security keys, the security key ID helps you determine which security key to enter.</p>

5.7.9 Verify the Security Key

Use the command in the following table to verify the security key.

Table 46: Verify the Security Key ID

Convention	MegaCli -VerifySecurityKey -SecurityKey ssssssssss -aN
Description	Verifies that the security key is the correct one for the self-encrypted disk. -VerifySecurityKey: Verifies the security key. -SecurityKey ssssssssss: Enters the new security key. The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.

5.8 Controller Property-Related Options

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

5.8.1 Display Controller Properties

Use the command in the following table to display parameters for the selected controllers.

Table 47: Controller Parameters

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

5.8.2 Display Number of Controllers Supported

Use the command in the following table to display the number of controllers supported on the system.

Table 48: Number of Controllers Supported

Convention	MegaCli -AdpCount
Description	Displays the number of controllers supported on the system and returns the number to the operating system.

5.8.3 Enable or Disable Automatic Rebuild

Use the command in the following table to turn automatic rebuild on or off for the selected controllers. 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 starts when a drive that was part of the drive group is reinserted.

Table 49: Enable or Disable Automatic Rebuild

Convention	MegaCli -AdpAutoRbld -Enbl -Dsbl -Dsply -aN -a0,1,2 -aALL
Description	Enables or disables automatic rebuild on the selected controllers. The -Dsply option shows the status of the automatic rebuild state.

5.8.4 Flush Controller Cache

Use the command in the following table to flush the controller cache on the selected controllers. This option sends the contents of cache memory to the virtual drives. If the MegaRAID system must be powered down rapidly, you must flush the contents of the cache memory to preserve data integrity.

Table 50: Cache Flush on Selected Controller

Convention	MegaCli -AdpCacheFlush -aN -a0,1,2 -aALL
Description	Flushes the controller cache on the selected controllers.

5.8.5 Set Controller Properties

This command sets the properties on the selected controllers. 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 51](#) to display the list of properties you can set for the controllers.

Table 51: Set Controller Properties

Convention	<pre>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 NCQEnbl NCQDsbl {SSDSMARTCpyBkEnbl -val} {MaintainPdFailHistoryEnbl -val} {EnblSpinDownUnConfigDrvs -val} AutoEnhancedImportEnbl AutoEnhancedImportDsbl {-UseFDEOnlyEncrypt -val} {-PrCorrectUncfgdAreas -val} -aN -a0,1,2 -aALL</pre>
Description	<p>Sets the properties on the selected controllers. 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: Maximum 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 MB, 2 - 1 GB.</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: If firmware detects inconsistency, then CC is aborted.</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 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> <p>NCQEnbl: Enable the native command queueing.</p> <p>NCQDsbl: Disable the native command queueing.</p> <p>SSDSMARTCpyBkEnbl: Enable copyback operation on Self-Monitoring Analysis and Reporting Technology (SMART) errors on a CacheCade software. Copyback is initiated when the first SMART error occurs on a SSD that is part of a virtual drive.</p> <p>MaintainPdFailHistoryEnbl: Enable maintenance of the history of a failed drive.</p> <p>EnblSpinDownUnConfigDrvs: Enable spindown of unconfigured drives.</p> <p>AutoEnhancedImportEnbl: Enable the automatic enhanced import of foreign drives.</p> <p>AutoEnhancedImportDsbl: Disable the automatic enhanced import of foreign drives.</p> <p>UseFDEOnlyEncrypt: Use encryption on FDE drives only.</p> <p>PrCorrectUncfgdAreas:</p>

5.8.6 Display Specified Controller Properties

Use the command in the following table to display specified properties on the selected controllers.

Table 52: Display Specified Controller Properties

Convention	MegaCli -AdpGetProp CacheFlushInterval RebuildRate PatrolReadRate BgiRate CCRate ReconRate SpinupDriveCount SpinupDelay CoercionMode PredFailPollInterval ClusterEnable BatWarnDsbl EccBucketSize EccBucketLeakRate EccBucketCount AlarmDsply AbortCCOnError AutoDetectBackPlaneDsbl CopyBackDsbl LoadBalanceMode SMARTCpyBkEnbl SSDSMARTCpyBkEnbl MaintainPdFailHistoryEnbl EnblSpinDownUnConfigDrvs EnblSSDPatrolRead NCQDsply UseFDEOnlyEncrypt WBSupport AutoEnhancedImportDsbl PrCorrectUncfgdAreas DsblSpinDownUnConfigDrvs ENABLEEGHSP ENABLEEUG ENABLEESMARTER -aN -a0,1,2 -aALL
Description	<p>Displays the properties on the selected controllers.</p> <p>EccBucketCount: Count of single-bit ECC errors currently in the bucket.</p> <p>WBSupport: Enables support for the Write Back option as the Write Policy.</p> <p>DsblSpinDownUnConfigDrvs: Disable spindown of unconfigured drives.</p> <p>ENABLEEGHSP: Enable global hot spare is three bits or adapter level for setting hot spare properties. Values: 0= Disable and 1: Enable.</p> <p>ENABLEEUG: Enable unconfigured good for emergency is three bits or adapter level for setting hot spare properties. Values: 0= Disable and 1: Enable</p> <p>ENABLEESMARTER: Emergency for SMARTer is three bits or adapter level for setting hot spare properties. Values: 0= Disable and 1: Enable</p> <p>See Table 51 for explanations of the other options.</p>

NOTE: The *tty* log can be saved at the controller level.

5.8.7 Set Factory Defaults

Use the command in the following table to set the factory defaults on the selected controllers.

Table 53: Set Factory Defaults

Convention	MegaCli -AdpFacDefSet -aN -a0,1,2 -aALL
Description	Sets the factory defaults on the selected controllers.

5.8.8 Set SAS Address

Use the command in the following table to set the SAS address on the selected controllers.

Table 54: Set SAS Address on Controller

Convention	MegaCli -AdpSetSASA str[0-64] -aN
Description	Sets the controller's SAS address. This string must be a 64-digit hexadecimal number.

5.8.9 Set Time and Date on Controller

Use the command in the following table to set the time and date on the selected controllers.

Table 55: Set Time and Date on Controller

Convention	MegaCli -AdpSetTime <i>yyyymmdd HH:mm:ss</i> -aN -a0,1,2 -aALL
Description	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.

5.8.10 Display Time and Date on Controller

Use the command in the following table to display the time and date on the selected controllers.

Table 56: Display Time and Date on Controller

Convention	MegaCli -AdpGetTime -aN
Description	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.

5.8.11 Get Connector Mode

Use the command in the following table to display which ports are enabled (Internal/External, 0/1) on the MegaRAID SAS 8888ELP RAID controller.

NOTE: This command is reserved strictly for the SAS 8888ELP RAID controller at this time. You must enable specific ports depending on how you intend to use the controller.

Table 57: Get Connector Mode

Convention	MegaCli -AdpGetConnectorMode -ConnectorN -Connector0,1 -ConnectorAll -aN -a0,1,2 -aALL
Description	Displays which ports are enabled (Internal/External, 0/1). For example, if internal port 0 is active, internal ports 0–3 are active. If external port 1 is active, external ports 4–7 are active.

5.8.12 Set Connector Mode

Use the command in the following table to set (enable) the connectors for the MegaRAID SAS 8888ELP RAID connectors that are listed in [Section 5.8.11, Get Connector Mode](#).

NOTE: This command is reserved strictly for the SAS 8888ELP RAID controller at this time. You must enable specific ports depending on how you intend to use the controller

Table 58: Set Connector Mode

Convention	MegaCli -AdpSetConnectorMode -Internal -External -Auto -ConnectorN -Connector0,1 -ConnectorAll -aN -a0,1,2 -aALL
Description	Sets (enables) the connectors listed in the GetConnectorMode command. For example, to enable internal ports 4-7 on controller 0, run the following command: MegaCli -AdpSetConnectorMode -Internal -Connector1 -a0

5.9 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.

5.9.1 Set Patrol Read Options

Use the command in the following table on the selected controllers to set the patrol read options.

Table 59: Set Patrol Read Options

Convention	MegaCli -AdpPR -Dsbl EnblAuto EnblMan Start Stop -Suspend -Resume Info {-SetStartTime yyyyymmdd hh} maxConcurrentPD -aN -a0,1,2 -aALL
Description	Sets the patrol read options on a single controller, multiple controllers, or all controllers: <ul style="list-style-type: none"> -Dsbl: Disables the patrol read for the selected controllers. -EnblAuto: Enables the patrol read automatically for the selected controllers. This means the patrol read will start automatically after the controller initialization is complete. -EnblMan: Enables the patrol read manually for the selected controllers. This means that the patrol read does not start automatically; it has to be started manually by selecting the Start command. -Start: Starts the patrol read for the selected controllers. -Stop: Stops the patrol read for the selected controllers. -Suspend: Suspends the patrol read. -Resume: Resumes a suspended partol read from the point that the patrol read was suspended. -Info: Displays the following patrol read information for the selected controllers: <ul style="list-style-type: none"> • Patrol read operation mode • Patrol read execution delay value • Patrol read status -SetStartTime yyyyymmdd hh: Set the start time for the patrol read in year/month/day format. -maxConcurrentPD: Sets the maximum number of concurrent drives on which the patrol read runs.

5.9.2 Set Patrol Read Delay Interval

Use the command in the following table on the selected controllers to set the time between patrol read iterations.

Table 60: Set Patrol Read Delay Interval

Convention	MegaCli -AdpPRSetDelay -Val -aN -a0,1,2 -aALL
Description	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.

5.10 BIOS-Related Properties

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

5.10.1 Set or Display Bootable Virtual Drive ID

Use the command in the following table 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 does not load if the boot sector is incorrect.

Table 61: Bootable Virtual Drive ID

Convention	MegaCli -AdpBootDrive {-Set -Lx -physdrv[E0:S0] } -Get -aN -a0,1,2 -aALL
Description	Sets or displays the bootable virtual drive ID: -Set -Lx -physdrv[E0:S0]: Sets the virtual drive as bootable so that during the next reboot, the BIOS looks for a boot sector in the specified virtual drive. Identifies the physical drive in the virtual drive, by enclosure and slot, to use to boot. -Get: Displays the bootable virtual drive ID.

5.10.2 Select BIOS Status Options

Use the command in the following table to set the options for the BIOS status.

Table 62: Options for BIOS Status

Convention	MegaCli -AdpBIOS -Enbl -Dsbl SOE BE EnblAutoSelectBootLd DsblAutoSelectBootLd -Dsply -aN -a0,1,2 -aALL
Description	Sets the BIOS options. The following are the settings you can select on a single controller, multiple controllers, or all controllers: -Enbl, -Dsbl: Enables or disables the BIOS status on selected controllers. -SOE: Stops on BIOS errors during POST for selected controllers. When set to -SOE, the BIOS stops in case of a problem with the configuration. This setting allows you to enter the configuration utility to resolve the problem. This setting is available only when you enable the BIOS status. -BE: Bypasses BIOS errors during the POST. This value is available only when you enable the BIOS status. -EnblAutoSelectBootLd DsblAutoSelectBootLd: Enable or disable automatic selection of the boot virtual drive. -Dsply: Displays the BIOS status on selected controllers.

5.11 Battery Backup Unit-Related Properties

5.11.1 Display BBU Information

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

Use the command in the following table to display complete information about the BBU for the selected controllers.

Table 63: Display BBU Information

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

5.11.2 Display BBU Status Information

Use the command in the following table to display complete information about the status of the BBU, such as temperature and voltage, for the selected controllers.

Table 64: Display BBU Status Information

Convention	MegaCli -AdpBbuCmd -GetBbuStatus -aN -a0,1,2 -aALL
Description	<p>Displays complete information about the BBU status, such as the temperature and voltage. The information displays in the following formats:</p> <p>BBU Status for Adapter: xx Battery Type: XXXXXX(string) Voltage: xx mV Current: xx mA Temperature: xx C° Firmware Status: xx Battery state: xx</p> <p>Gas Gauge Status: 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>For iBBU: 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>For BBU: Relative State of Charge: xx Charger Status: xx Remaining Capacity: xx mAh Full Charge Capacity: mAh isSOHGood: Yes/No</p>

5.11.3 Display BBU Capacity

Use the command in the following table to display the BBU capacity for the selected controllers.

Table 65: Display BBU Capacity Information

Convention	MegaCli -AdpBbuCmd -GetBbuCapacityInfo -aN -a0,1,2 -aALL
Description	Displays BBU capacity information. The information displays in the following formats: 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%

5.11.4 Display BBU Design Parameters

Use the command in the following table to display BBU design parameters for the selected controllers.

Table 66: Display BBU Design Parameters

Convention	MegaCli -AdpBbuCmd -GetBbuDesignInfo -aN -a0,1,2 -aALL
Description	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)

5.11.5 Display Current BBU Properties

Use the command in the following table to display the current BBU properties for the selected controllers.

Table 67: Display Current BBU Properties

Convention	MegaCli -AdpBbuCmd -GetBbuProperties -aN -a0,1,2 -aALL
Description	<p>Displays current properties of the BBU. The information displays in the following formats:</p> <p>BBU Properties for Adapter: x</p> <p>Auto Learn Period: xxx Sec</p> <p>Next Learn Time: xxxx Sec</p> <p>Learn Delay Interval=<value>: Value in hours, not greater than 168 hours (7 days)</p> <p>Auto-Learn Mode=<value>: Value can be 0, 1, or 2.</p>
	<p>NOTE: If the battery type is IBBU08, then the BBU mode is displayed as a part of GetBbuProperties.</p>
	<p>NOTE: When the value in the Auto Learn Mode is set from 1 (Disabled) to 0 (Enabled), the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. The Learn Delay Interval field and the Next Learn Time field will not be updated until the battery relearn is completed. Once the relearning cycle is completed, the value in the Next Learn Time field and in the Learn Delay Interval field will display the new time (in seconds) of the next battery learning cycle</p>

5.11.6 Start BBU Learning Cycle

Use the command in the following table to start the BBU learning cycle on the selected controllers. 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 68: Start BBU Learning Cycle

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

5.11.7 Place Battery in Low-Power Storage Mode

Use the command in the following table to place the battery into Low-Power Storage mode on the selected controllers. This mode saves battery power consumption.

Table 69: Place Battery in Low-Power Storage Mode

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

5.11.8 Set BBU Properties

Use the command in the following table to set the BBU properties on the selected controllers after reading from the file.

Table 70: Set BBU Properties

Convention	MegaCli -AdpBbuCmd -SetBbuProperties -f<fileName> -aN -a0,1,2 -aALL
Description	Sets the BBU properties on the selected controllers 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.
	NOTE: You can change only two of these parameters: learnDelayInterval and autoLearnMode.
	NOTE: If the battery type is IBBU08, then in autoLearnmode, the WarnViaEvent is not supported.

5.12 Options for Displaying Logs Kept at the Firmware Level

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

5.12.1 Event Log Management

Use the command in the following table to manage the event entries in the event log for the selected controllers.

Table 71: Event Log Management

Convention	MegaCli -AdpEventLog -GetEventLogInfo -aN -GetEvents {-info -warning -critical -fatal} GetSinceShutdown {-info -warning -critical -fatal} GetSinceReboot {-info -warning -critical -fatal} IncludeDeleted {-info -warning -critical -fatal} {GetLatest <number> {-info -warning -critical -fatal} } -f <filename> Clear -aN -a0,1,2 -aALL {GetCCIncon} -f <filename> -LX -L0,2,5... -LALL -aN -a0,1,2 -aALL
Description	Manages event log entries. The following are the settings you can select on a single controller, multiple controllers, or all controllers: -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. -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. <i>Start_entry</i> specifies the initial event log entry when displaying the log. -GetSinceShutdown: Displays all of the events since last controller shutdown. -GetSinceReboot: Displays all of the events since last controller reboot. -IncludeDeleted: Displays all events, including deleted events. -GetLatest: Displays the latest number of events, if any exist. The event data will be writtent to the file in reverse order. -Clear: Clears the event log for the selected controllers. -GetCCIncon: Displays the events relating to inconsistent data found during a consistency check.

NOTE: `-AdpEventLogInfo` does not support the file option in PCLl.

5.12.2 Set BBU Terminal Logging

Use the command in the following table to set the BBU terminal logging for the selected controllers.

Table 72: Set BBU Terminal Logging

Convention	MegaCli -FwTermLog -Bbuoff -BbuoffTemp -Bbuon -BbuGet -aN -a0,1,2 -aALL
Description	<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"> -Bbuoff: While storing the TTY log in DRAM, BBU is not used for buffering. In the case of power loss, this log is lost.. -BbuoffTemp: TTY BBU buffering will be set to off only for this boot. -Bbuon: While storing the TTY log in DRAM, BBU is used for buffering. The Log is preserved even in the case of power loss. -BbuGet: This parameter gives the current BBU state, i.e., if BBU is on or off for TTY history.

5.13 Configuration-Related Options

You can specify the drives by using the Enclosure ID:Slot ID for SAS controllers. This option assumes that all drives are connected to the controller through an enclosure. If the drives are not connected to an enclosure, they are assumed to be connected to Enclosure 0. In this case no slot exists, so you can use the `pdlist` command to get the slot equivalent number. (This option 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.

5.13.1 Create a RAID Drive Group from All Unconfigured-Good Drives

Use the command in the following table to create one RAID drive group out of all of the Unconfigured Good drives, and a hot spare, if desired. This command 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 drives of different capacities exist, the largest drive is used to make the hot spare.

NOTE: The firmware supports only 32 drives per drive group. If more than 32 Unconfigured Good drives exist, MegaCli cannot configure any of the drives, and the command aborts.

Table 73: Create a Drive Group from All of the Unconfigured Drives

Convention	MegaCli -CfgLDAdd -RX[E0:S0,E1:S1,...] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-szXXX [-szYYY ...]] [-strpszM] [-Hsp[E0:S0,...]] [-AfterLdX] -Force [FDE CtrlBased]
Description	<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, 1, 5, or 6. 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 the write policy.</p> <p>-NORA (No Read Ahead), RA (Read Ahead), ADRA (Adaptive Read Ahead): Selects the read policy.</p> <p>-Direct, -Cached: Selects the cache policy.</p> <p>-CachedBadBBU NoCachedBadBBU: Specifies whether to use write cache when the BBU is bad.</p> <p>Hsp: Specifies the drive with which to make the hot spare.</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> <hr/> <p>NOTE: Previously, -szXXX expressed capacity in MB, but now you can enter the capacity in your choice of units. For example, to create a virtual drive of 10 GB, enter the size as sz10GB. If you do not enter a unit, MB is the default unit.</p>

5.13.2 Add RAID 0, 1, 5, or 6 Configuration

Use the command in the following table 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 5.13.3, Add RAID 10, 50, or 60 Configuration](#).

Table 74: Add RAID 0, 1, 5, or 6 Configuration

Convention	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] [-Force] -aN
Description	<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, because 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 the write policy.</p> <p>-NORA (No Read Ahead), RA (Read Ahead), ADRA (Adaptive Read Ahead): Selects the read policy.</p> <p>-Cached, -Direct: Selects the cache policy.</p> <p>[{CachedBadBBU NoCachedBadBBU }]: Specifies : 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 MB. 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.</p> <p>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.</p> <p>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 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 KB.</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.</p> <p>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>

5.13.3 Add RAID 10, 50, or 60 Configuration

Use the command in the following table 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 5.13.2, Add RAID 0, 1, 5, or 6 Configuration](#).

Table 75: Add RAID 10, 50, or 60 Configuration

Convention	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] -Force [FDE CtrlBased] -aN -a0,1,2 -aALL
Description	<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 8 spans with a maximum of 32 drives in each span. (Some factors, such as the type of controller, limit the number of drives you can use.) RAID 10 requires an even number of drives, because data from one drive is mirrored to the other drive in each RAID 1 drive group. You can have an even number or odd number of spans.</p> <p>Multiple drive groups are specified using the <code>-ArrayX[E0:S0, ...]</code> option. (Note that X 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 <code>-afterLdX</code> option is specified, but the controller does not support slicing in the spanned drive groups.</p>
	<p>NOTE: Previously, <code>-szXXX</code> expressed capacity in MB, but now you can enter the capacity in your choice of units. For example, to create a virtual drive of 10 GB, enter the size as <code>sz10GB</code>. If you do not enter a unit, MB is the default unit.</p>

5.13.4 Clear the Existing Configuration

Use the command in the following table to clear the existing configuration on the selected controllers.

Table 76: Clear Existing Configuration

Convention	MegaCli -CfgClr -aN -a0,1,2 -aALL
Description	Clears the existing configuration.

5.13.5 Save the Configuration on the Controller

Use the command in the following table to save the configuration for the selected controllers to the given file name.

Table 77: Save Configuration on the Controller

Convention	MegaCli -CfgSave -f FileName -aN
Description	Saves the configuration for the selected controllers to the given filename.

5.13.6 Restore the Configuration Data from File

Use the command in the following table to read the configuration from the file and load it on the selected controllers. You can restore the read/write properties and RAID configuration using hot spares.

Table 78: Restore Configuration Data from File

Convention	MegaCli -CfgRestore -f FileName -aN
Description	<p>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:</p> <ul style="list-style-type: none"> • MegaCli does not validate the setup when restoring the RAID configuration. • 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 fails if the same device IDs of the drives are not present.

5.13.7 Manage Foreign Configuration Information

Use the command in the following table to manage configurations from other controllers, called *foreign configurations*, for the selected controllers. 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 the -Preview option before you import a foreign configuration.

Table 79: Manage Foreign Configuration Information

Convention	MegaCli -CfgForeign -Scan [-SecurityKey ssssssssss] -Dsply [x] [-SecurityKey ssssssssss] -Preview [x] [-SecurityKey ssssssssss] -Import [x] [-SecurityKey ssssssssss] -Clear [x] [-SecurityKey ssssssssss] -aN -a0,1,2 -aALL
Description	<p>Manages foreign configurations. The options for this command follow:</p> <ul style="list-style-type: none"> -Scan: Scans and displays available foreign configurations. -SecurityKey: This key is based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. If the security key is unavailable, user data is irretrievably lost. You must be careful to never lose the security key. -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.

5.13.8 Delete Specified Virtual Drives

Use the command in the following table to delete one, multiple, or all virtual drives on the selected controllers.

Table 80: Delete Specified Virtual Drives

Convention	MegaCli -CfgLDDel -Lx -L0,1,2 -La11 -aN -a0,1,2 -aALL
Description	Deletes the specified virtual drives on the selected controllers. You can delete one virtual drive, multiple virtual drives, or all of the selected virtual drives on selected controllers.

5.13.9 Display the Free Space

Use the command in the following table to display the free space that is available to use for configuration on the selected controllers.

Table 81: Display Free Space

Convention	MegaCli -CfgFreeSpaceInfo -aN -a0,1,2 -aALL
Description	Displays all of the free space available for configuration on the selected controllers. 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.

5.14 Virtual Drive-Related Options

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

5.14.1 Display Virtual Drive Information

Use the command in the following table to display virtual drive information for the selected controllers.

Table 82: Display Virtual Drive Information

Convention	MegaCli -LDInfo -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Displays information about the virtual drives on the selected controllers. 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.

5.14.2 Change the Virtual Drive Cache and Access Parameters

Use the command in the following table to change the cache policy and access policy for the virtual drives on the selected controllers.

Table 83: Change Virtual Drive Cache and Access Parameters

Convention	MegaCli -LDSetProp WT WB [-Immediate] RA NORA ADRA -Cached Direct CachedBadBBU NoCachedBadBBU } -RW RO Blocked {-Name nameString} -EnDskCache DisDskCache -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Allows you to change the following virtual drive parameters: -WT (Write through), WB (Write back): Selects write policy. -Immediate: Indicates that the changes take place immediately. -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.

5.14.3 Display the Virtual Drive Cache and Access Parameters

Use the command in the following table to display cache and access parameters for the virtual drives on the selected controllers.

Table 84: Display Virtual Drive Cache and Access Parameters

Convention	MegaCli -LDGetProp -Cache -Access -Name -DskCache -Lx -L0, 1, 2 -Lall -aN -a0, 1, 2 -aALL
Description	Displays the cache and access policies of the virtual drives: -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.

5.14.4 Manage Virtual Drives Initialization

Use the command in the following table to manage initialization of the virtual drives on the selected controllers.

Table 85: Manage Virtual Drive Initialization

Convention	MegaCli -LDInit {-Start [Fast Full]} -Abort -ShowProg -ProgDsply -Lx -L0, 1, 2 -Lall -aN -a0, 1, 2 -aALL
Description	Allows you to select the following actions for virtual drive initialization: -Start: Starts the initialization (writing 0s) on the virtual drives 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 drives. -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.

5.14.5 Manage a Consistency Check

Use the command in the following table to manage a data consistency check (CC) on the virtual drives for the selected controllers.

Table 86: Manage Consistency Check

Convention	MegaCli -LDCC -Start -Abort -Suspend -Resume -ShowProg -ProgDsply -Lx -L0, 1, 2 -Lall -aN -a0, 1, 2 -aALL
Description	Allows you to select the following actions for a data CC: -Start: Starts a CC on the virtual drives, then displays the progress (optional) and time remaining. -Abort: Aborts an ongoing CC on the virtual drives. -Suspend: Suspends the CC. -Resume: Resumes a CC from the point where the CC was suspended. -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.

5.14.6 Schedule a Consistency Check

Use the command in the following table to schedule a consistency check (CC) on the virtual drives for the selected controllers. There are options to set the mode, change the CC start time, set the delay time and display of the CC info.

Table 87: Schedule Consistency Check

Convention	MegaCli -AdpCcSched -Dsbl -Info { -ModeConc -ModeSeq [-ExcludeLd -LN -L0,1,2] [-SetStartTime yyyyymmdd hh] [-SetDelay val] } -aN -a0,1,2 -aALL
Description	<p>Schedules check consistency on the virtual drive of the selected adapter.</p> <p>Dsbl: Disables a scheduled CC for the given adapters.</p> <p>Info: Gets information about a scheduled CC for the given adapters.</p> <p>ModeConc: The scheduled CC on all of the virtual drives runs concurrently for the given adapters.</p> <p>ModeSeq: The scheduled CC on all of the virtual drives runs sequentially for the given adapters</p> <p>ExcludeLd: Specify the virtual drive numbers not included in the scheduled CC. The new list will overwrite the existing list stored on the controller. This is optional.</p> <p>StartTime: Sets the next start time. The date is in the format of yyyyymmdd in decimal digits and followed by a decimal number for the hour between 0 ~ 23 inclusively. This is optional.</p> <p>SetDelay: Sets the execution delay between executions for the given adapters. This is optional.</p> <p>Values: The value is the length of delay in hours. A value of 0 means continuous execution.</p>

5.14.7 Manage a Background Initialization

Use the command in the following table to enable, disable, or suspend background initialization (BGI), as well as display initialization progress on the selected controllers.

Table 88: Manage Background Initialization

Convention	MegaCli -LDBI -Enbl -Dsbl -getSetting -abort -Suspend -Resume -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Manages background initialization options. The following are the background initialization settings you can select on a single controller, multiple controllers, or all controllers:</p> <p>-Enbl, -Dsbl: Enables or disables the background initialization on the selected controllers.</p> <p>-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.</p> <p>-Abort: Aborts an ongoing background initializations.</p> <p>-Suspend: Suspends the background initializations.</p> <p>-Resume: Resumes a background initializations from the point where the background initializations was suspended.</p> <p>-ShowProg: Displays the current progress value.</p> <p>- GetSetting: Displays current background initialization setting (Enabled or Disabled).</p>

5.14.8 Perform a Virtual Drive Reconstruction

Use the command in the following table to perform a reconstruction of the virtual drives on the selected controllers.

Table 89: Virtual Drive Reconstruction

Convention	MegaCli -LDRecon {-Start -Rx [Add Rmv PhysDrv[E0:S0,E1:S1,...]] } -ShowProg -ProgDsply -Lx -aN
Description	<p>Controls and manages virtual drive reconstruction. The following are the virtual drive reconstruction settings you can select on a single controller:</p> <ul style="list-style-type: none"> -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.

5.14.9 Display Information about Virtual Drives and Drives

Use the command in the following table to display information about the virtual drives and drives for the selected controllers, such as the number of virtual drives, RAID level, and drive capacity.

Table 90: Display Virtual Drive and Drive Information

Convention	MegaCli -LDPDInfo -aN -a0,1,2 -aALL
Description	<p>Displays information about the present virtual drives and drives on the selected controllers. The command displays the following information.</p> <ul style="list-style-type: none"> • The number of virtual drives. • The RAID level of the virtual drives. • The device world-wide name. • The device firmware level. • The device write-cache setting. • The device negotiated transfer speed (link speed) for each active or passive port. • The device's disk group membership. • An indication if the device has flagged a SMART alerts. • The status of each physical port on the physical device (if it is active, passive or disabled). • The firmware version of the device. • The new PD state (UnConfigured - Shielded, Hot Spare - shielded, Configured - shielded). • The Shield Counter value. • The last shield diagnostics completion time. • The drive capacity information, which includes raw capacity, coerced capacity, uncoerced capacity, drive temperature, enclosure position and SAS address. • For SATA devices, it indicates if NCQ is supported/enabled or disabled.

5.14.10 Display the Bad Block Table

Use the command in the following table to check for bad block entries of virtual disks on the selected adapter.

Table 91: Display Virtual Drive and Drive Information

Convention	MegaCLI -GetBbtEntries -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Displays information on the bad block entries of virtual disks on the selected adapters.

5.14.11 Display the Number of Virtual Drives

Use the command in the following table to display the number of virtual drives attached to the controller.

Table 92: Display Number of Virtual Drives

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

5.14.12 Clear the LDBBM Table Entries

Use the command in the following table to clear the LDBBM table entries.

Table 93: Clear the LDBBM Table Entries

Convention	MegaCli -LDBBMClr -Lx -L0,1,2,... -Lall -aN -a0,1,2 -aALL
Description	Clears the LDBBM table entries for the virtual drives on the selected adapters.

5.14.13 Display the List of Virtual Drives with Preserved Cache

Use the command in the following table to display the list of virtual drives that have preserved cache. Preserved cache is cache that remains in the controller cache after a drive goes offline or missing and that has not been saved to a drive yet. You can reboot and manage the preserved cache.

Table 94: Display the List of Virtual Drives with Preserved Cache

Convention	MegaCli -GetPreservedCacheList -aN -a0,1,2 -aALL
Description	Displays the list of virtual drives that have preserved cache.

5.14.14 Discard the Preserved Cache of a Virtual Drives

Use the command in the following table to discard the preserved cache of a virtual drives.

Table 95: Discard the Preserved Cache of a Virtual Drives

Convention	MegaCli -DiscardPreservedCache -Lx -L0,1,2 -Lall -force -aN -a0,1,2 -aALL
Description	Discard the preserved cache of the virtual drives.

5.14.15 Expand a Virtual Drive

Use the command in the following table to expand a virtual drive.

Table 96: Discard the Preserved Cache of a Virtual Drives

Convention	MegaCli -LdExpansion -pN -dontExpandArray -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Expands the virtual drive within the existing array or beyond the size of the existing array if you replace the drives with larger drives.</p> <p>-pN: Denotes the percentage of the array to use to expand the virtual drive. N ranges from 0 to 100 percent. For example, -p30 indicates expansion up to 30 percent of available array size.</p> <p>-dontExpandArray: Expand a virtual drive within the array, even when there is room to expand the array.</p> <p>For example, you have created a 5-GB RAID 1 virtual drive with two 30-GB drives. The array size is 30 GB and the virtual drive size is 5 GB. If you replace the two 30-GB drives with two 60-GB drives, the array size is still 30 GB (because of previous configuration). You have two options:</p> <ol style="list-style-type: none"> 1. Expand the virtual drive within the array. Use the -dontExpandArra option to expand the virtual drive up to 30 GB. 2. Expand the virtual drive beyond the existing array size Use the -pN option to expand the virtual drive beyond 30 GB and up to 60 GB (the size of the replacement drives).

5.15 Drive-Related Options

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

5.15.1 Display Drive Information

Use the command in the following table to display information about the drives on the selected controllers.

Table 97: Display Drive Information

Convention	MegaCli -PDInfo -PhysDrv [E0:S0,E1:S1 . . .] -aN -a0,1,2 -aALL
Description	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, inquiry data, device world-wide name, device firmware level, device write-cache setting, device negotiated transfer speed (link speed) for each active or passive port, device's disk group membership, if the device has flagged a S.M.A.R.T. alert, the status of each physical port on the physical device (if it is active, passive or disabled) and firmware version of the device. For SAS devices, this includes additional information, such as the SAS address of the drive. For SAS expanders, this command includes additional information, such as the number of devices connected to the expander. -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives about which to provide information.

5.15.2 Set the Drive State to Online

Use the command in the following table 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 98: Set Drive State to Online

Convention	MegaCli -PDOnline -PhysDrv [E0:S0,E1:S1 . . .] -aN -a0,1,2 -aALL
Description	Changes the drive state to Online. -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.15.3 Set the Drive State to Offline

Use the command in the following table to set the state of a drive to *Offline*. In the offline state, the virtual drive is not available to the RAID controller.

Table 99: Set Drive State to Offline

Convention	MegaCli -PDOffline -PhysDrv [E0:S0,E1:S1 . . .] -aN -a0,1,2 -aALL
Description	Changes the drive state to Offline. -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.15.4 Change the Drive State to Unconfigured-Good

Use the command in the following table to change the state of a drive from Unconfigured-Bad to Unconfigured-Good.

Table 100: Change Drive State to Unconfigured Good

Convention	MegaCli -PDMakeGood -PhysDrv[E0:S0,E1:S1...] [-Force] -aN -a0,1,2 -aALL
Description	Changes the drive state to Unconfigured-Good. -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives. Force: Force the drive to the Unconfigured-Good state.

5.15.5 Change the Drive State

Use the command in the following table to change the drive state, as it relates to hot spares, and to associate the drive to an enclosure and to a drive group for the selected controllers.

Table 101: Change Drive State

Convention	MegaCli -PDHSP {-Set [{-Dedicated -ArrayN -Array0,1...}] [-EnclAffinity] [-nonRevertible] } -Rmv -PhysDrv[E0:S0,E1:S1,...] -aN -a0,1,2 -aALL
Description	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 dedicated hot spare for the enclosure. -Array0: Dedicates the hot spare to a specific drive group number N. -EnclAffinity: Associates the hot spare to a selected enclosure. -Rmv: Changes the drive state toReady (removes the hot spare). -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives. You can get the list of arrays by using the CLI command CfgDsply. In the results of the CfgDsply command, the number associated with DISK GROUPS is the array number.

5.15.6 Manage a Drive Initialization

Use the command in the following table to manage a drive initialization on the selected controllers.

Table 102: Drive Initialization

Convention	MegaCli -PDClear -Start -Stop -ShowProg -ProgDsply -PhysDrv [E0:S0, E1:S1 . . .] -aN -a0, 1, 2 -aALL
Description	<p>Manages initialization or displays initialization progress on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"> -Start: Starts initialization on the selected drives. -Stop: Stops an ongoing initialization on the selected drives. -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.

5.15.7 Rebuild a Drive

Use the command in the following table 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 re-creating the data that was stored on the drive before it failed.

Table 103: Rebuild a Drive

Convention	MegaCli -PDRbld -Start -Stop -Suspend -Resume -ShowProg -ProgDsply -PhysDrv [E0:S0, E1:S1 . . .] -aN -a0, 1, 2 -aALL
Description	<p>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:</p> <ul style="list-style-type: none"> -Start: Starts a rebuild on the selected drives and displays the rebuild progress (optional). -Stop: Stops an ongoing rebuild on the selected drives. -Suspend: Suspends the rebuild. -Resume: Resumes the rebuild from the point that the rebuild was suspended. -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. -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.15.8 Locate the Drives and Activate LED

Use the command in the following table to locate the drives for the selected controllers and activate the Drive Activity LED.

Table 104: Locate Drive and Activate LED

Convention	MegaCli -PDLocate -PhysDrv [E0:S0,E1:S1...] -aN -a0,1,2 -aALL
Description	Locates the drives for the selected controllers and activates the Drive Activity LED. -Physdrv [E0:S0, ...]: Specifies the physical drive enclosure and the slots for the drives.

5.15.9 Mark the Configured Drive as Missing

Use the command in the following table to mark the configured drive as missing for the selected controllers.

Table 105: Mark Configured Drive as Missing

Convention	MegaCli -PdMarkMissing -PhysDrv [E0:S0,E1:S1...] -aN -a0,1,2 -aALL
Description	Marks the offline drive as missing for the selected controllers. -Physdrv [E0:S0, ...]: Specifies the physical drive enclosure and the slots for the drives.

NOTE: -PdMarkMissing works only on offline drives. If you want to make a configured drive as missing, first mark it as offline, and then mark it as missing. When PdReplaceMissing is run, the drive becomes offline, and rebuild does not start automatically. You have to start it explicitly.

Follow these steps to replace the PD or retrieve the PD:

1. pdgetmissing (This command reports the array and the row number needed for the next command).
2. pdreplacemissing (input the array and row number here).
3. pdonline.

5.15.10 Display the Drives in Missing Status

Use the command in the following table to mark the configured drive as missing for the selected controllers.

Table 106: Display Drives in Missing Status

Convention	MegaCli -PDGetMissing -aN -a0,1,2 -aALL												
Description	Displays the drives in missing status. The format follows. <table border="0"> <tr> <td>No</td> <td>Row</td> <td>Column</td> <td>SizeExpected(MB)</td> </tr> <tr> <td>0</td> <td>x</td> <td>y</td> <td>zzzzzzzz</td> </tr> <tr> <td colspan="4">...</td> </tr> </table> <p>Where x is the index to the drive groups, y is the index to the drive in that drive group, and zzzzzz is the minimum capacity of the drive that can be used as a replacement.</p>	No	Row	Column	SizeExpected(MB)	0	x	y	zzzzzzzz	...			
No	Row	Column	SizeExpected(MB)										
0	x	y	zzzzzzzz										
...													

5.15.11 Replace the Configured Drives and Start an Automatic Rebuild

Use the command in the following table to replace configured drives and start an automatic rebuild of the drive for the selected controllers.

Table 107: Replace Configured Drives and Start Automatic Rebuild

Convention	MegaCli -PDReplaceMissing -PhysDrv[E0:S0,E1:S1...] -ArrayX -RowY -aN
Description	Replaces the configured drives that are identified as missing and then starts an automatic rebuild. -Physdrv[E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.15.12 Prepare the Unconfigured Drive for Removal

Use the command in the following table to prepare the unconfigured drives for removal from the selected controllers.

Table 108: Prepare Unconfigured Drives for Removal

Convention	MegaCli -PDPrpRmv [-Undo] - PhysDrv[E0:S0,E1:S1...] -aN -a0,1,2 -aALL
Description	Prepares unconfigured drives for removal. The firmware spins down this drive. The drive state is set to Unaffiliated, which marks it as offline even though it is not a part of a configuration. -Undo: This option undoes this operation. If you select undo, the firmware marks this drive as Unconfigured Good. -Physdrv[E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.15.13 Display Total Number of Drives

Use the command in the following table to display the total number of drives attached to an controller. Drives can be attached directly or through enclosures.

Table 109: Display Number of Drives Attached to an Controller

Convention	MegaCli -PDGetNum -aN -a0,1,2 -aALL
Description	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.

5.15.14 Display List of Physical Devices

Use the command in the following table to display a list of the physical devices connected to the selected controllers.

Table 110: Display List of Physical Devices Attached to Controllers

Convention	MegaCli -PDList -aN -a0,1.. -aAll
Description	<p>Displays information about all drives and other devices connected to the selected controllers. This command displays the following information</p> <ul style="list-style-type: none"> • The information such as the drive type, • The capacity (if a drive), • The serial number, • The drive temperature, enclosure position, • The device world-wide name, • The device firmware level, • The device write-cache setting, • The device negotiated transfer speed (link speed) for each active or passive port, device's disk group membership, • If the device has flagged a S.M.A.R.T. alert, • The status of each physical port on the physical device (if it is active, passive or disabled) and firmware version of the device. <p>For SAS devices, it includes additional information such as the SAS address of the device.</p> <p>For SAS expanders, it includes additional information such as the number of drives connected to the expander.</p> <p>For SATA devices, it indicates whether NCQ is supported/enabled or disabled.</p>

5.15.15 Download Firmware to the Physical Devices

Use the command in the following table to download firmware to the physical devices connected to the selected controllers.

Table 111: Download Firmware to the Physical Devices

Convention	MegaCli -PdFwDownload [offline] {[-SataBridge] -PhysDrv[E0:S0] } {-EncdevId[devId]} -f <filename> -aN -a0,1,2 -aALL (new one)
Description	<p>Flashes the firmware with the file specified at the command line. The firmware files used to flash a physical device can be of any format. The CLI utility assumes that you provide a valid firmware image, and it flashes the same. The physical device needs to do error checking.</p> <p>-SataBridge: Allows you to download the SATA bridge firmware in online mode.</p> <p>-Physdrv[E0:S0]: Specifies the physical drive enclosure and the slots for the drives. Flashes the firmware to only one PD.</p> <p>-EncdevId[devId1, devId2, . . .]: Specifies the enclosure device ID. See Section 5.16.1, Display Enclosure Information for more enclosure information.</p>

NOTE: The PdFwDownload command does not support LSI SAS1 expander firmware upgrade.

5.15.16 Configure All Free Drives into a RAID 0, 1, 5, or 6 Configuration for a Specific Controller

Use the command in the following table to download firmware to the physical devices connected to the selected controllers.

Table 112: Configure All Free Drives into a RAID 0, 1, 5 or 6 Configuration for a Specific Controller

Convention	<pre>MegaCli -CfgAllFreeDrv -rX [-SATAOnly] [-SpanCount XXX] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-strpszM] [-HspCount XX [-HspType -Dedicated -EnclAffinity -nonRevertible]] [FDE CtrlBased] -aN</pre>
Description	<p>Adds all of the unconfigured physical drives to a RAID level 0, 1, 5, or 6 configuration on a specified controller. Even if no configuration is present, you have the option to write the configuration to the controller.</p> <p><code>rX[E0:S0, . . .]</code>: Specifies the RAID level and the physical drive enclosure/slot numbers to construct a disk group.</p> <p><code>WT</code> (Write Through), <code>WB</code> (Write Back): Selects the write policy.</p> <p><code>NORA</code> (No Read Ahead), <code>RA</code> (Read Ahead), <code>ADRA</code> (Adaptive Read Ahead): Selects the read policy.</p> <p><code>Cached</code>, <code>-Direct</code>: Selects the cache policy.</p> <p><code>[{CachedBadBBU NoCachedBadBBU}]</code>: Specifies whether to use write cache when the BBU is bad.</p> <p><code>szXXXXXXXX</code>: Specifies the size for the virtual disk, where <code>XXXX</code> is a decimal number of MB. However, the actual size of the virtual drive might be smaller, because the driver requires the number of blocks from the physical drives in each virtual drive to be aligned to the stripe size.</p> <p>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 fails if the remaining size of the array is too small to configure the virtual drive with the specified size. This option can also be used to create a configuration on the free space available in the array.</p> <p><code>strpszM</code>: Specifies the stripe size, where the stripe size values are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, or 1024 MB.</p> <p><code>Hsp[E5:S5, . . .]</code>: Creates hot spares when you create the configuration. The new hot spares are dedicated to the virtual drive used to create the configuration. This option does not allow you to create global hot spares. To create global hot spares, you must use the <code>-PdHsp</code> command with the proper subcommands.</p> <p>You can also use this option to create a configuration on the free space available in the virtual drive.</p> <p><code>AfterLdX</code>: 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 disk is already used for configuration.</p> <p><code>FDE CtrlBased</code>: If the controller supports the security feature, this option enables FDE/controller-based encryption on the virtual disk.</p>

5.15.17 Set the Mapping Mode of the Drives to the Selected Controllers

Use the command in the following table to set the mapping mode of the physical devices connected to the selected controllers.

Table 113: Set the Mapping Mode of the Drive to the Selected Controller

Convention	MegaCli -DirectPdMapping -Enbl -Dsbl -Dsply -aN -a0,1,2 -aALL
Description	Sets the mapping mode of the drives connected to the specified controllers. Enbl: Enables the direct physical drive mapping mode. Dsbl: Disables the direct physical drive mapping mode. Dsply: Displays the current state of the direct physical drive mapping.

5.15.18 Secure Erase for Virtual Drives and Physical Drives

Use the command in the following table to the perform the secure erase operation on a virtual drive or a physical drive.

The command in this section performs a secure erase. It performs a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns.

Table 114: Secure Erase for Virtual Drives and Physical Drives

Convention	MegaCli -SecureErase Start -EraseType [[-ErasePattern A] [-ErasePattern B]] Stop ShowProg ProgDsply -PhysDrv [E0:S0,E1:S1,...] -Lx -L0,1,2 -LALL -aN -a0,1,2 -aALL
Description	The SecureErase command performs a series of write operations to a drive that overwrite every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security.

5.15.19 Perform the Copyback Operation on the Selected Drive

Use the command in the following table to the perform the copyback operation on the selected drive.

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).

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.

Table 115: Perform the Copyback Operation on the Selected Drive

Convention	MegaCli -PDCpyBk -Start -Stop -Suspend -Resume -ShowProg -ProgDsply -PhysDrv [E0 : S0] -aN -a0, 1, 2 -aALL
Description	Performs the copyback operation on the selected physical drive. Start: Initializes the copyback operation on the selected drive. Stop: Stops the copyback operation on the selected drive. Suspend: Suspends the copyback operation. Resume: Resumes a copyback operation from the point that the copyback operation was suspended. ShowProg: Displays a snapshot of the ongoing copyback operation. ProgDsply: Allows you to view the ongoing copyback operation. The routine continues to display progress until at least one copyback is completed or a key is pressed. -Physdrv [E0 : S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

5.16 Enclosure-Related Options

5.16.1 Display Enclosure Information

The commands in this section are used for enclosures.

Use the command in the following table to display enclosure information for selected controllers.

Table 116: Display Enclosure Information

Convention	MegaCli -EncInfo -aN -a0, 1, 2 -aALL
Description	Displays information about the enclosure for the selected controllers. The following properties are displayed. <ul style="list-style-type: none"> • The enclosure type. • The enclosure serial number. • The ESM serial number. • The firmware version installed. • The chassis status. • The chassis temperature or threshold status (e.g. normal, high, etc.). It is the same as enclosure temperature from enclosure status. • The fan status (e.g., normal, missing, etc.). • The power supply count. • The power supply status for each installed power supply. • The VPD field replaceable unit (FRU) part number. • The enclosure zoning mode. • The enclosure vendor identifier.

5.16.2 Display Enclosure Status

Use the command in the following table to display the status of the enclosure for selected controllers.

Table 117: Display Enclosure Status

Convention	MegaCli -EncStatus -aN -a0,1,2 -aALL
Description	Displays the status of the enclosure for the selected controllers.

5.16.3 Upgrading the Firmware without Restarting

Use the command in the following table to upgrade the firmware in the enclosure without restarting the enclosure.

Convention	PdFwDownload [offline][ForceActivate] {[-SataBridge] -PhysDrv[0:1] } {-EncdevId[devId1]} -f <filename> -aN -a0,1,2 -aALL
Description	The ForceActivate suboption enables you to upgrade the firmware in the enclosure and activate the upgraded firmware without restarting the enclosure.

5.17 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.

5.17.1 Flash the Firmware with the ROM File

Use the command in the following table to flash the firmware with the ROM file specified at the command line for the selected controllers.

Table 118: Flash Firmware with ROM File

Convention	MegaCli -AdpFwFlash -f filename [-NoSigChk] [-NoVerChk]-aN -a0,1,2 -aALL
Description	<p>Flashes the firmware with the ROM file specified at the command line.</p> <p>The -NoSigChk 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 -NoVerChk, 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"> • Select which controller to flash after the controllers are detected. • Flash the firmware on all present controllers. • The option generates an XML output data.

5.17.2 Flash the Firmware in Mode 0 with the ROM File

Use the command in the following table to flash the firmware in Mode 0 with the ROM file specified at the command line for the selected controllers. This command is only supported for MS-DOS.

Table 119: Flash Firmware in Mode 0 with ROM File

Convention	MegaCli -AdpM0Flash -f filename
Description	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 function, is to flash the firmware on all present controllers which are compatible with the image.

5.18 SAS Topology

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

Use the command in the following table to display the PHY connection information for physical PHY M on the selected controllers. 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 120: Display PHY Connection Information

Convention	MegaCli -PHYInfo -phyM -aN -a0,1,2 -aALL
Description	Displays PHY connection information for physical PHY M on the controllers.

5.19 Diagnostic-Related Options

5.19.1 Start Controller Diagnostics

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

Use the command in the following table to start the controller diagnostic for a set amount of time.

Table 121: Start Diagnostics Setting

Convention	MegaCli -AdpDiag [val] -aN -a0,1,2 -aALL
Description	Sets the amount of time for the controller diagnostic to run. Val: Indicates the time in seconds for the controller diagnostic to run.

5.19.2 Perform a Full Stroke Seek Test

Use the command in the following table to perform a full stroke seek. This command is only supported for MS-DOS.

Table 122: Start Full Stroke Seek Test

Convention	MegaCli - FullStrSeekTest a0
Description	This CLI function is used for testing the server's power supply capability to withstand all drives doing a full stroke seek. Upon receipt of the CLI command the firmware ceases normal operation, seeks all the drives to cylinder 0, and waits 5 seconds. After 5 seconds, it seeks for all drives to the last cylinder. The Full Stroke Seek Test is a continuous operation that runs until power cycle

5.19.3 Start Battery Test

Use the command in the following table to start the battery test. This command requires a system reboot.

Table 123: Start Battery Test

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

5.20 Recovery (Snapshot)-Related Options

The commands in this section are used to perform actions with the Recovery advanced software, also known as Snapshot (LSIP200038104).

The Recovery feature uses Snapshot technology to offer a simplified way to recover lost data and provides protection for any volume, including the boot volume. You can use the Recovery feature to take snapshots of a volume at designated point-in-time (PIT) and restore the volume or files from those points in case data is deleted, whether accidentally or maliciously. MegaRAID Recovery supports up to 8 snapshots of PITs for each volume.

5.20.1 Enable the Snapshot Feature

Use the command in the following table to enable the snapshot feature on a selected virtual drive..

Table 124: Enable the Snapshot Feature

Convention	MegaCli -Snapshot -Enbl -szXXX SnapshotRepositoryLD N [-AutoSnapshot] [AutoDeleteOldestSnapshot] -Lx -aN -a0,1,2 -aALL
Description	Enables the snapshot on the source virtual drive for the corresponding snapshot target virtual drive. -szXXX: Specifies the size in MB on for the virtual drive, where XXX is a decimal number of MB. SnapshotRepositoryLD N: Specifies the repository LD number. -AutoSnapshot: Optional parameter, if specified, enables the AutoSnapshot for the source virtual drive. -AutoDeleteOldestSnapshot: Optional parameter, if specified, enables the AutoDeleteOldestSnapshot for the source virtual drive. -Lx: x specifies the source LD number on which to enable snapshot.

5.20.2 Disable the Snapshot Feature

Use the command in the following table to enable the snapshot feature on a selected virtual drive.

Table 125: Disable the Snapshot Feature

Convention	MegaCli -Snapshot -Dsb1 -Lx -aN -a0,1,2 -aALL
Description	Disables the snapshot on the source virtual drive. -Lx: x specifies the source LD number on which to disable snapshot.

5.20.3 Take a Snapshot of a Volume

Use the command in the following table to take a snapshot of a volume at designated point-in-time .

Table 126: Take Snapshot of Volume

Convention	MegaCli -Snapshot -TakeSnapshot [-snapshotName name] [-CreateView [-ViewName view_name] [-RW RO Blocked] [-szXXX]] -LN -L0,1,2 -aN -a0,1,2 -aALL
Description	Takes a snapshot of a volume at designated point-in-time. -snapshotName name: (Optional) If specified, the snapshot is created with the name you enter for it. -CreateView: (Optional) If specified, this option creates a view for the snapshot. A view contains the content from the point-in-time [PIT] when the snapshot was made. -ViewName view_name: (Optional) Specifies the name of the view you created. -RW RO Blocked: (Optional) Specifies the access policy of the view. -szXXX: Specifies the size of the view in MB where XXX is a decimal number. -LN: N specifies the source LD number for the command.

5.20.4 Set the Snapshot Properties

Use the command in the following table to set the snapshot properties..

Table 127: Set the Snapshot Properties

Convention	MegaCli -Snapshot -SetProp {-AutoSnapshot -val} {-AutoDeleteOldestSnapshot -val} -Lx -aN -a0,1,2 -aALL
Description	Sets the Snapshot properties, such as AutoSnapshot and AutoDeleteOldestSnapshot. -AutoSnapshot: If the value is 0, this command disables the AutoSnapshot feature on source virtual drive. If the value is 1, it enables the AutoSnapshot feature on source virtual drive. -AutoDeleteOldestSnapshot: If the value is 0, this command disables the AutoDeleteOldestSnapshot feature on the source virtual drive. If the value is 1, it enables the AutoDeleteOldestSnapshot feature on the source virtual drive. -Lx: x specifies the source LD number for the command.

5.20.5 Delete a Snapshot

Use the command in the following table to delete a snapshot.

Table 128: Delete a Snapshot

Convention	MegaCli -Snapshot -DeleteSnapshot [SnapshotTime yyyyymmdd hh:mm:ss -all] [-force -y] -LN -L0,1,2 -aN -a0,1,2 -aALL
Description	<p>Deletes the snapshot and the associated view if <code>-Force</code> or <code>-Y</code> is specified.</p> <p><code>-SnapshotTime yyyyymmdd hh:mm:ss</code>: (Optional) If used, this action deletes the snapshot with the time stamp that is specified in command line, if it is the oldest PiT.</p> <p><code>-force</code>: If specified, this action deletes the snapshot even if it has the view associated with it.</p> <p><code>-y</code>: If specified, this action deletes the snapshot even if it has the view associated with it.</p> <p><code>-LN: N</code> specifies the source LD number for the command.</p> <p><code>-L0, 1, 2</code>: Specifies the command is for LDs 0, 1, and 2. You can select more than one LD.</p>

5.20.6 Create a View

Use the command in the following table to create a view. A view contains the content from the PiT when the snapshot was made.

Table 129: Create a View

Convention	MegaCli -Snapshot -CreateView -SnapshotTime yyyyymmdd hh:mm:ss [-viewName NameString] [-RW RO Blocked] [-szXXX] -Lx -aN -a0,1,2 -aALL \n", appNameP);
Description	<p>Creates the view on a particular snapshot.</p> <p><code>-SnapshotTime yyyyymmdd hh:mm:ss</code>: Creates the view on the snapshot with the time stamp <code>yyyyymmdd hh:mm:ss</code></p> <p><code>-viewName NameString</code>: (Optional) Specifies the name of the view.</p> <p><code>-RW RO Blocked</code>: (Optional) Specifies the access policy of the view.</p> <p><code>-szXXX</code>: (Optional) Specifies the size of the view in MB where <code>XXX</code> is a decimal number.</p> <p><code>-Lx: x</code> specifies the source LD number for the command.</p>

5.20.7 Delete a View

Use the command in the following table to a view..

Table 130: Delete a View

Convention	MegaCli -Snapshot -DeleteView [-SnapshotTime yyyyymmdd hh:mm:ss] -Lx -aN -a0,1,2 -aALL
Description	<p>Deletes the view.</p> <p><code>-SnapshotTime yyyyymmdd hh:mm:ss</code>: (Optional) If specified, this action deletes the view on the snapshot with the time stamp <code>yyyyymmdd hh:mm:ss</code>.</p> <p><code>-Lx: x</code> specifies the source LD number for the command.</p>

5.20.8 Roll back to an Older Snapshot

Use the command in the following table to roll the virtual drive back to an older snapshot..

Table 131: Rollback to an Old Snapshot

Convention	MegaCli -Snapshot -Rollback -SnapshotTime <i>yyyymmdd hh:mm:ss</i> [-Force -Y] -Lx -aN -a0,1,2 -aALL
Description	<p>Rolls back the virtual drive to an old snapshot. The Rollback option is supported by Preboot MegaCli not by the OS level MegaCli.</p> <p>-SnapshotTime <i>yyyymmdd hh:mm:ss</i>: Specifies the snapshot with the time stamp <i>yyyymmdd hh:mm:ss</i> to which it has to roll back.</p> <p>-Force: If specified, this option overrides the warning message and causes a rollback to an older snapshot.</p> <p>-Y: If specified, this option overrides the warning message and causes a rollback to an older snapshot.</p> <p>-Lx: <i>x</i> specifies the source LD number for the command.</p>

5.20.9 Display Snapshot and View Information

Use the command in the following table to display information about the snapshot and the view..

Table 132: Display Snapshot and View Information

Convention	MegaCli -Snapshot -Info [-SnapshotTime <i>yyyymmdd hh:mm:ss</i> -ViewTime <i>yyyymmdd hh:mm:ss</i>] -Lx -aN -a0,1,2 -aALL
Description	<p>Displays snapshot and view information for the source virtual drive. If the virtual drive is a repository virtual drive, it displays the LD information, the number of source virtual drives mapped and their target IDs and the number of holes.</p> <p>-SnapshotTime <i>yyyymmdd hh:mm:ss</i>: (Optional) If specified, this displays the snapshot information for the snapshot with the time stamp <i>yyyymmdd hh:mm:ss</i>.</p> <p>-ViewTime <i>yyyymmdd hh:mm:ss</i>: (Optional) If specified, this displays the view information for the view with the time stamp <i>yyyymmdd hh:mm:ss</i> and the associated snapshot information.</p> <p>-Lx: <i>x</i> specifies the source LD number for the command.</p>

5.20.10 Clean the Recoverable Free Space on the Drives in a Virtual Drive

Use the command in the following table to clean the recoverable free space on the drives in a snapshot repository virtual drive. The free space is unused space on the drives in a virtual drive.

Table 133: Clean the Recoverable Free Space on the Drives in a Virtual Drive

Convention	MegaCli -Snapshot -Clean -Lx -aN -a0,1,2 -aALL
Description	<p>Cleans the recoverable free space on the drives in a snapshot repository virtual drive.</p> <p>-Lx: <i>x</i> specifies the LD number for the command. The LD must be a repository virtual drive.</p>

5.20.11 Display the Information for a Specific View

Use the command in the following table to display the information for a specific view if you specify the view target ID.

Table 134: Display the Information for a Specific View

Convention	MegaCli -Snapshot -GetViewInfo [-ViewTargetId N] -aN -a0,1,2 -aALL
Description	Displays the view information about a particular view if you specify the -ViewTargetId. Otherwise, it displays the information about all of the views. -ViewTargetId N: (Optional) If specified, this displays the information about the view with the specified target ID.

5.20.12 Enable the Snapshot Scheduler

The snapshot scheduler in the MegaRAID Storage Manager software helps you automate the creation of point-in-time (PiT) on one or more virtual discs. You can schedule the snapshot as monthly, weekly, daily, or hourly. The scheduler does not support biweekly, alternate days, etc.

NOTE: The MegaRAID Storage Manager software or OEM applications must be running to flush file system buffers and take snapshots when snapshot schedule expires. The system does not support snapshot creation through BMC (sideband) and snapshot on volumes, which are used for virtual machine creation under virtualized environment. To flush the file system buffers, you should interface with the hypervisor.

5.21 FastPath-Related Options

The command in this section displays information about the FastPath option.

MegaRAID FastPath is a high-performance IO accelerator for CacheCade 2.0 SSD Read Caching software drive groups connected to a MegaRAID controller card. CacheCade 2.0 SSD Read Caching software has a read performance advantage over HDDs and uses less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 6 Gb/s MegaRAID SATA+SAS controller.

The FastPath feature supports full optimization of CacheCade 2.0 SSD Read Caching software and hard disk drive (HDD) virtual drive groups to deliver a three-fold improvement in read and write IOPS compared to MegaRAID controllers not using FastPath technology. Also, FastPath software is faster and more cost-effective than current flash-based adapter card solutions.

5.22 Dimmer Switch-Related Options

5.22.1 Display Selected Adapter Properties

The following tables display command information about the Dimmer Switch option.

Use the command in the following table to display the selected adapter properties.

Table 135: Displays the Selected Adapter Properties

Convention	MegaCLI -AdpGetProp CacheFlushInterval RebuildRate PatrolReadRate BgiRate CCRate ReconRate SpinupDriveCount SpinupDelay CoercionMode ClusterEnable PredFailPollInterval BatWarnDsbl EccBucketSize EccBucketLeakRate EccBucketCount AbortCCOnError AlarmDsply SMARTCpyBkEnbl SSDSMARTCpyBkEnbl NCQDsply MaintainPdFailHistoryEnbl RstrHotSpareOnInsert DisableOCR EnableJBOD DsblCacheBypass BootWithPinne dCache AutoEnhancedImportDsply AutoDetectBackPlaneDs bl EnblSpinDownUnConfigDrvs SpinDownTime DefaultSnap shotSpace DefaultViewSpace AutoSnapshotSpace CopyBac kDsbl LoadBalanceMode UseFDEOnlyEncrypt UseDiskActivityForLocate DefaultLdPSPolicy DisableLdPsInterval DisableLdPsTime SpinUpEncDrvCn SpinUpEncDelay -aN -a0, 1, 2 -aALL
Description	<p>Displays selected adapter properties.</p> <p>The possible settings follow:</p> <p>DefaultLdPSPolicy: Default LD power savings policy.</p> <p>DisableLdPsInterval: LD power savings are disabled for yy hours beginning at disableLdPsTime.</p> <p>DisableLdPsTime: LD power savings shall be disabled at xx minutes from 12:00 am.</p> <p>SpinUpEncDrvCnt: Maximum number of drives within an enclosure to spin up at one time.</p> <p>SpinUpEncDelay: Number of seconds to delay among spinup groups within an enclosure.</p>

5.22.2 Sets the Properties on the Selected Adapter

Use the command in the following table to set the properties on the selected adapter.

Table 136: Sets the Properties on the Selected Adapter

Convention	<pre>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} {SSDSMARTCpyBkEnbl -val} NCQEnbl NCQDsbl {MaintainPdFailHistoryEnbl -val} {RstrHotSpareOnInsert -val} {EnblSpinDownUnConfigDrvs -val} {DisableOCR -val} {BootWithPinnedCache -val} AutoEnhancedImportEnbl AutoEnhancedImportDsbl {CopyBackDsbl -val} {AutoDetectBackPlaneDsbl -val} {LoadBalanceMode -val} {UseFDEOnlyEncrypt -val} {DsblSpinDownHsp -val} {SpinDownTime -val} {EnableJBOD -val} {DsblCacheBypass -val} {useDiskActivityForLocate -val} {SpinUpEncDrvCnt -val} {SpinUpEncDelay -val} {-ENABLEEGHSP -val} {-ENABLEEUG -val } {ENABLEESMARTER -val} -aN -a0,1,2 -aALL</pre>
Description	<p>Sets the properties on the selected adapters.</p> <p>The possible settings follow:</p> <p>SpinUpEncDrvCnt: Max number of drives within an enclosure to spin up at one time. Values: 0 to 255.</p> <p>SpinUpEncDelay: Number of seconds to delay among spinup groups within an enclosure. Values: 0 to 255.</p> <p>ENABLEEGHSP: Enable global hot spare is 3 bits or adapter level for setting hot spare properties. Values: 0= Disable and 1= Enable.</p> <p>ENABLEEUG: Enable unconfigured good for emergency is 3 bits or adapter level for setting hot spare properties. Values: 0= Disable and 1= Enable</p> <p>ENABLEESMARTER: Emergency for SMARTer is 3 bits or adapter level for setting hot spare properties. Values: 0= Disable and 1:=Enable</p>

5.22.3 Displays the Power-Saving Level on the Virtual Disk

Use the command in the following table to display the power-saving level on the virtual disk.

Table 137: Displays the Power Saving Level on the Virtual Disk

Convention	<pre>MegaCLI -LDSetPowerPolicy -Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL</pre>
Description	Sets the power-saving level on the virtual disk.

5.22.4 Displays about Adding a RAID Level to a Specified Adapter

Use the command in the following table to add a RAID level to a specified adapter.

Table 138: Adds a RAID Level to a Specified Adapter

Convention	MegaCLI -CfgLdAdd -rX[E0:S0,E1:S1,...] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-szXXX [-szYYY ...]] [-strpszM] [-Hsp[E0:S0,...]] [-AfterLdX] [-Force] [FDE CtrlBased] [-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] -aN
Description	<p>-CfgLdAdd: Adds a RAID level 0, 1, 5, or 6 to a specified adapter. Even if no configuration is present, you have the option to write the configuration to the adapter.</p> <p>The possible parameters follow:</p> <p>Rx[E0:S0, . . .]: Specifies the RAID level and the physical drive enclosure/slot numbers to construct a disk array.</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 size for the virtual disk, where XXXX is a decimal number of Mbytes. However, the actual size of the virtual disk may be smaller, because the driver requires the number of blocks from the physical drives in each virtual disk to be aligned to the stripe size. If multiple size options are specified, CT will configure the virtual disks in the order of the options entered in the command line. The configuration of a particular virtual disk fails if the remaining size of the array is too small to configure the virtual disk with the specified size. This option can also be used to create a configuration on the free space available in the array.</p> <p>strpszM: Specifies the strip size, where the strip size values are 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, 256 MB, 512 MB, or 1024 MB.</p> <p>Hsp[E5:S5, . . .]: Creates hot spares when you create the configuration. The new hot spares will be dedicated to the virtual disk 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 the proper subcommands. You can also use this option to create a configuration on the free space available in the virtual disk.</p> <p>AfterLdX: This command is optional. By default, the application uses the first free slot available in the virtual disk. This option is valid only if the virtual disk is already used for configuration.</p> <p>Force: This option forces the creation of virtual disk in situations where the application finds that it is convenient to create the virtual disk only with user's consent.</p> <p>FDE CtrlBased: If the controller supports the security feature, this option enables FDE/control-based encryption on the virtual disk.</p> <p>[-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching]: If the controller supports power savings on virtual disk, these options specify the possible levels of power savings that can be applied on a virtual disk.</p>

5.22.5 Create a RAID Level

Use the command in the following table to create a RAID level 10, 50, 60 (spanned configuration).

Table 139: Creates a RAID Level 10, 50, 60 (spanned) Configuration

Convention	<pre>MegaCLI -CfgSpanAdd -r10 -Array0[E0:S0,E1:S1] -Array1[E0:S0,E1:S1] [-ArrayX[E0:S0,E1:S1] ...] [WT WB] [NORA RA ADRA] [Direct Cached] [CacheBadBBU NoCacheBadBBU] [-strpszM] [-szXXX[-szYYY ...]] [-AfterLdX] [-Force] [FDE CtrlBased] [-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] -aN MegaCLI -CfgSpanAdd -r50 -Array0[E0:S0,E1:S1,E2:S2,...] -Array1[E0:S0,E1:S1,E2:S2,...] [-ArrayX[E0:S0,E1:S1,E2:S2,...] ...] [WT WB] [NORA RA ADRA] [Direct Cached] [CacheBadBBU NoCacheBadBBU] [-strpszM] [-szXXX[-szYYY ...]] [-AfterLdX] [-Force] [FDE CtrlBased] [-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] -aN</pre>
Description	<p>-CfgSpanAdd: Creates a RAID level 10, 50, or 60 (spanned) configuration from the specified arrays. Even if no configuration is present, you must use this option to write the configuration to the adapter.</p> <p>The possible parameters are:</p> <p>Rx: Specifies the RAID Level.</p> <p>ArrayX[E0:S0, ...]: Specifies the Array and the physical drive enclosure/slot numbers to construct a disk array.</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>Cache-Direct: Selects cache policy.</p> <p>[{CacheBadBBU NoCacheBadBBU }]: Specifies whether to use write cache when the BBU is bad.</p> <p>szXXXXXXXX: Specifies the size for the virtual disk, where XXXX is a decimal number of Mbytes. However, the actual size of the virtual disk may be smaller, because the driver requires the number of blocks from the physical drives in each virtual disk to be aligned to the stripe size. If multiple size options are specified, CT will configure the virtual disks in the order of the options entered in the command line. The configuration of a particular virtual disk fails if the remaining size of the array is too small to configure the virtual disk with the specified size. This option can also be used to create a configuration on the free space available in the array.</p> <p>strpszM: Specifies the strip size, where the strip size values are 8, 16, 32, 64, 128, 256, 512, or 1024 MB.</p> <p>AfterLdX: This command is optional. By default, the application uses the first free slot available in the virtual disk. This option is valid only if the virtual disk is already used for configuration.</p> <p>Force: This option forces the creation of virtual disk in situations where the application finds that it is convenient to create the virtual disk only with user's consent.</p> <p>FDE CtrlBased: If the controller supports the security feature, this option enables FDE/Ctrl based encryption on the virtual disk.</p> <p>[-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] : If the controller supports power savings on virtual disk, these options specify the possible levels of power savings that can be applied on a virtual disk.</p>

5.22.6 Add the Unconfigured Drive to a Specified Adapter

Use the command in the following table to add the unconfigured drives to an adapter.

Table 140: Adds the Unconfigured Physical Drive to RAID Level 0, 1, 5, 6 to a Specified Adapter

Convention	<pre>MegaCLI -CfgAllFreeDrv -rX [-SATAOnly] [-SpanCount XXX] [WT WB] [NORA RA ADRA] [Direct Cached] [CachedBadBBU NoCachedBadBBU] [-strpszM] [-HspCount XX [-HspType -Dedicated -EnclAffinity -nonRevertible]] [FDE CtrlBased] [-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] -aN</pre>
Description	<p>Adds all the unconfigured physical drives to RAID level 0, 1, 5, or 6 configuration to a specified adapter. Even if no configuration is present, you have the option to write the configuration to the adapter.</p> <p>The possible parameters are:</p> <p>Rx [E0 : S0 , . . .]: Specifies the RAID level and the physical drive enclosure/slot numbers to construct a disk array.</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 size for the virtual disk, where XXXX is a decimal number of Mbytes. However, the actual size of the virtual disk may be smaller, because the driver requires the number of blocks from the physical drives in each virtual disk to be aligned to the strip size. If multiple size options are specified, CT will configure the virtual disks in the order of the options entered in the command line. The configuration of a particular virtual disk fails if the remaining size of the array is too small to configure the virtual disk with the specified size. This option can also be used to create a configuration on the free space available in the array.</p> <p>strpszM: Specifies the strip size, where the strip size values are 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, 256 MB, 512 MB, or 1024 MB.</p> <p>Hsp [E5 : S5 , . . .]: Creates hot spares when you create the configuration. The new hot spares will be dedicated to the virtual disk 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 the proper subcommands. You can also use this option to create a configuration on the free space available in the virtual disk.</p> <p>AfterLdX: This command is optional. By default, the application uses the first free slot available in the virtual disk. This option is valid only if the virtual disk is already used for configuration.</p> <p>FDE CtrlBased: If controller support security feature, this option enables FDE/Ctrl based encryption on virtual disk.</p> <p>[-Default -Automatic -None -MaximumWithCaching -MaximumWithoutCaching] : If the controller supports power savings on virtual disk, these options specify the possible levels of power savings that can be applied on a virtual disk.</p>

5.22.7 Displays the Cache and Access Policies

Use the command in the following table to display the cache and access policies of the virtual disks.

Table 141: Displays the Cache and Access Policies of the Virtual Disks

Convention	MegaCLI -LDGetProp -Cache -Access -Name -DskCache -PSPolicy -Lx -L0,1,2 -LALL -aN -a0,1,2 -aALL
Description	Displays the cache and access policies of the virtual disks The possible parameters follow: 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 physical disk cache policy. PSPolicy: Displays the default and current power savings policy of the virtual disk.

5.23 Performance Monitoring Options

The commands in this section are used to monitor the performance of the system.

5.23.1 Starting Performance Data Collection

Use this command to start the collection of performance data for the time interval (in minutes) specified by you. Once the specified time has elapsed, the performance data collection stops.

Table 142: Starting Performance Data Collection

Convention	MegaCli -perfmon -start -interval <val> -aN
Description	-perfmon: Specifies collection of performance data. The possible parameters are: -start: Starts the performance data collection. -interval: the time interval (in minutes) that you specified for which you want the performance data collection. The default value is 5 minutes.

5.23.2 Stopping Performance Data Collection

Use this command if the performance data collection is already running and you want to stop the collection of performance data.

Table 143: Stopping Performance Data Collection

Convention	MegaCli -perfmon -stop -aN
Description	-perfmon: Specifies collection of performance data. The possible parameters are: -stop: Stops the performance data collection.

5.23.3 Saving Performance Data

Use this command to save the performance data collection in a file.

Table 144: Saving Performance Data

Convention	MegaCli -perfmon -getresults -f <Filename> -aN
Description	-perfmon: Specifies collection of performance data. The possible parameters are: -getresults: Specifies saving the performance data collection. -f: Specifies the file name in which the performance data is to be saved. The format of the file is CSV.

5.24 Miscellaneous Commands

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

5.24.1 Display the Version

Use this command to display the CLI version, the version of the device driver, the version of the Unified extended firmware interface (UEFI) device driver, the firmware versions for the attached physical device, and the enclosure.

Convention	MegaCli -Version -Cli Ctrl Driver Pd Uefi aN (Uefi works only for EFI)
Description	Displays the firmware versions and other code levels installed on the controller, the CLI version, the version of the device driver, the version of the UEFI device driver, the firmware versions for the attached physical device, and enclosure in a list as location information, model string, and the firmware version.

5.24.2 Display the MegaCLI Version

Use the command in the following table to display the version number of the MegaCLI utility.

Table 145: Display MegaCLI Version

Convention	MegaCli -v
Description	Displays the version number of the MegaCLI utility.

5.24.3 Display Help for MegaCLI

Use the command in the following table to display help information for the MegaCLI utility.

Table 146: Display Help for MegaCLI

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

5.24.4 Display Summary Information

Use the command in the following table to show summary information for the MegaCLI utility.

Table 147: Display Help for the MegaCLI Utility

Convention	<code>MegaCli -ShowSummary [-f filename] -aN</code>
Description	Displays a summary of the system information, controller information, the drive information, the virtual drive information, and the enclosure information.

Chapter 6

MegaRAID Storage Manager Overview and Installation

This chapter provides a brief overview of the MegaRAID® Storage Manager software and explains how to install it on the supported operating systems.

6.1 Overview

The MegaRAID Storage Manager software enables you to configure, monitor, and maintain storage configurations on LSI SAS controllers. The MegaRAID Storage Manager graphical user interface (GUI) makes it easy for you to create and manage storage configurations.

6.1.1 Creating Storage Configurations

The MegaRAID Storage Manager software enables you to easily configure the controllers, drives, and virtual drives on your workstation or on the server. The Configuration wizard greatly simplifies the process of creating drive groups and virtual drives. The wizard allows you to easily create new storage configurations and modify the configurations.

You can create configurations using the following modes:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you. This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize the creation of virtual drives. This option provides greater flexibility when creating virtual drives for your specific requirements because you can select the drives and the virtual drive settings when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

In addition, the Modify Drive Group wizard enables you to increase the capacity of a virtual drive and to change the RAID level of a drive group.

NOTE: The Modify Drive Group wizard was previously known as the Reconstruction wizard.

6.1.2 Monitoring Storage Devices

The MegaRAID Storage Manager software displays the status of controllers, virtual drives, and drives on the workstation or on the server that you are monitoring. The system errors and events are recorded in an event log file and are displayed on the dialog. Special device icons appear on the window to notify you of drive failures and other events that require immediate attention.

6.1.3 Maintaining Storage Configurations

You can use the MegaRAID Storage Manager software to perform system maintenance tasks, such as running patrol read operations, updating firmware, and running consistency checks on drive groups that support redundancy.

6.2 Hardware and Software Requirements

The hardware requirements for the MegaRAID Storage Manager software are as follows:

- PC-compatible computer with an IA-32 (32-bit) Intel Architecture processor or an EM64T (64-bit) processor; also compatible with SPARC V9 architecture-based systems
- Minimum 256 MB of system memory (512 MB recommended)
- A hard drive with at least 400 MB available free space; Solaris 10 SPARC occupies a minimum of 640 MB.

The supported operating systems for the MegaRAID Storage Manager software are as follows:

- Microsoft® Windows Server 2003, Microsoft Windows Server 2008, Microsoft Windows Server 2008R2, Microsoft Windows XP, Microsoft Windows Vista, and Microsoft Windows 7
- Red Hat Linux 3.0, 4.0, 5.0, and 6.0
- Solaris 10 x86, SPARC
- SUSE Linux/SLES 9, 10, and 11, with the latest updates and service packs
- VMware ESX 3.5, 4.0, and 4.1

Refer to your server documentation and to the operating system documentation for more information on hardware and operating system requirements.

You can use the MegaRAID Storage Manager software to remotely monitor the systems running the VMware ESXi (3.5 and above) operating system.

NOTE: Storelib libraries need the capability to be installed with more than one version. All the storelib libraries have been moved to a private location. Please do a clean un-installation and only then install the MegaRAID Software Manager to avoid any conflicts.

6.3 Prerequisites to Running MegaRAID Storage Manager Remote Administration

The MegaRAID Storage Manager software requires ports 3071 and 5571 to be open in order to function. Follow these steps to make sure these ports are open and to configure multicasting.

1. Configure the system with a valid Internet Protocol (IP) address.

Make sure there is no IP address conflict within the sub-network, and make sure that ports 3071 and 5571 are open and available for the MegaRAID Storage Manager framework communication.

2. Disable all security management and the firewall.
3. Configure multicasting.

Make sure Class D multicast IP addresses are registered (at least 229.111.112.12 should be registered for the MegaRAID Storage Manager software to work). If they are not registered, create a static route using the following command:

```
Route add 229.111.112.12 dev eth1
```

4. Install the MegaRAID Storage Manager software. If the MegaRAID Storage Manager software is already installed, restart the MegaRAID Storage Manager framework.

6.4 Installing MegaRAID Storage Manager

6.4.1 Prerequisite for MegaRAID Storage Manager Installation

This section explains how to install (or reinstall) the MegaRAID Storage Manager software on your workstation or on your server for the supported operating systems: Microsoft Windows, Red Hat Linux, and SUSE Linux.

The MegaRAID Storage Manager software installation script also installs the LSI SNMP agent, Red Hat Package Manager (RPM). The LSI SNMP agent application depends upon the standard SNMP-Util package.

Make sure that the SNMP-Util package is present in the system before you install the MegaRAID Storage Manager software.

The SNMP-Util package includes the RPM's net-snmp-libs, net-snmp-utils, and additional dependent RPM's. Make sure that these RPM's are installed from the operating system media before you install the MegaRAID Storage Manager software.

6.4.2 Installing MegaRAID Storage Manager Software on Microsoft Windows

To install the MegaRAID Storage Manager software on a system running the Microsoft Windows Server 2003, Microsoft Windows Server 2008, Microsoft Server 2008R2, Microsoft Windows XP, Microsoft Windows Vista, or Microsoft Windows 7 operating system, perform the following steps:

1. Insert the MegaRAID Storage Manager software installation CD in the CD-ROM drive.

If necessary, find and double-click the `setup.exe` file to start the installation program.

2. In the Welcome screen that appears, click **Next**.

If the MegaRAID Storage Manager software is already installed on this system, then an upgraded installation occurs.

3. Read and accept the user license in the following figure and click **Next**.



Figure 164: License Agreement

The Customer Information window appears, as shown in the following figure.

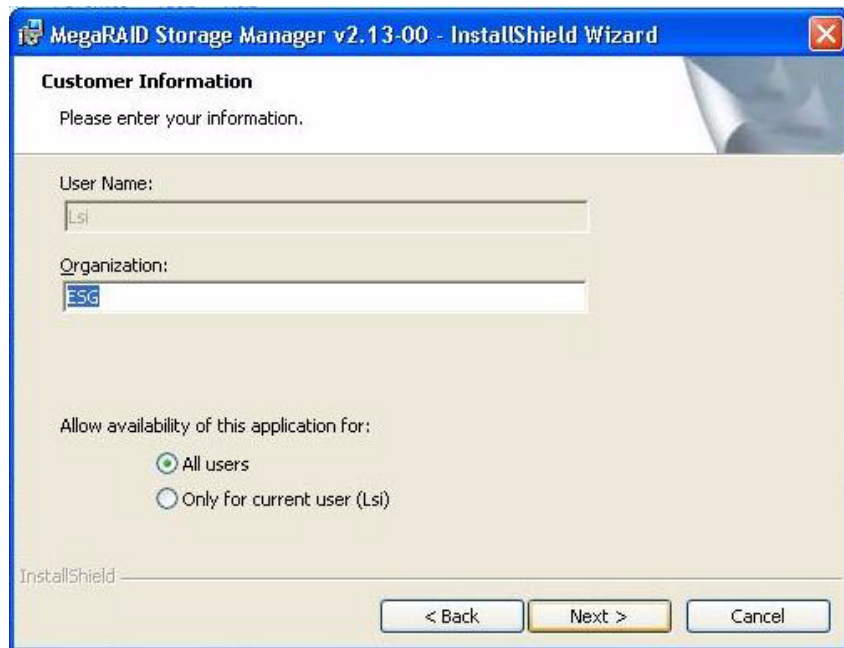


Figure 165: Customer Information Window

4. Enter your user name and organization name. In the bottom part of the screen, select an installation option:
 - If you select the **All users** radio button, any user with administrative privileges can use this version of the MegaRAID Storage Manager software to view or change storage configurations.
 - If you select the **Only for current user** radio button, the MegaRAID Storage Manager software shortcuts and associated icons are available only to the user with this user name.
5. Click **Next** to continue.

6. Accept the default destination folder, or click **Change** to select a different destination folder, as shown in the following figure.




Figure 166: Destination Folder Window

7. Click **Next** to continue.

The Setup Type window appears, as shown in the following figure.

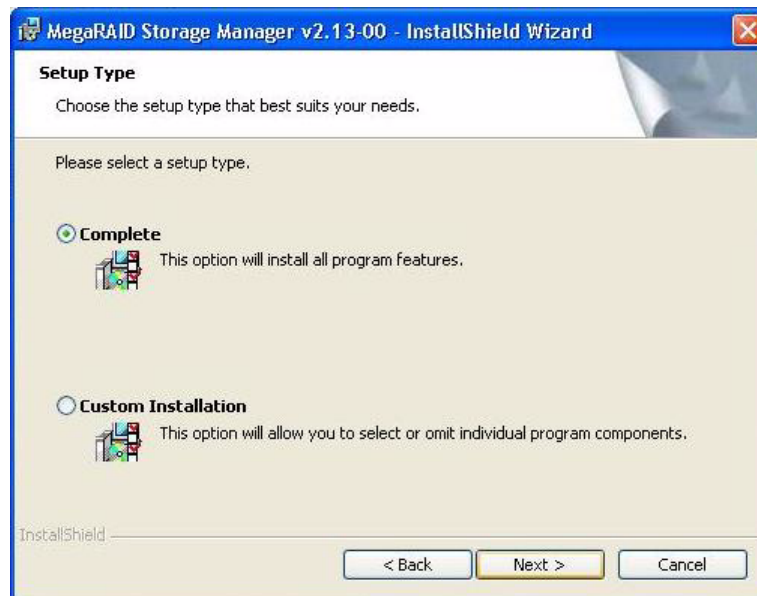


Figure 167: Setup Type Window

8. Select one of the setup options. The options are fully explained in the window text.
 - Select the **Complete** radio button if you are installing the MegaRAID Storage Manager software on a server.
 - Select the **Custom Installation** radio button if you want to select individual program components.
9. Click **Next** to continue.

If you selected **Custom Installation** as your setup option, the second Setup Type window appears, as shown in [Figure 168](#).

If you selected **Complete** as your setup option, the Installation wizard is ready to install the MegaRAID Storage Manager software. To begin installation, click **Install** on the next window that appears.

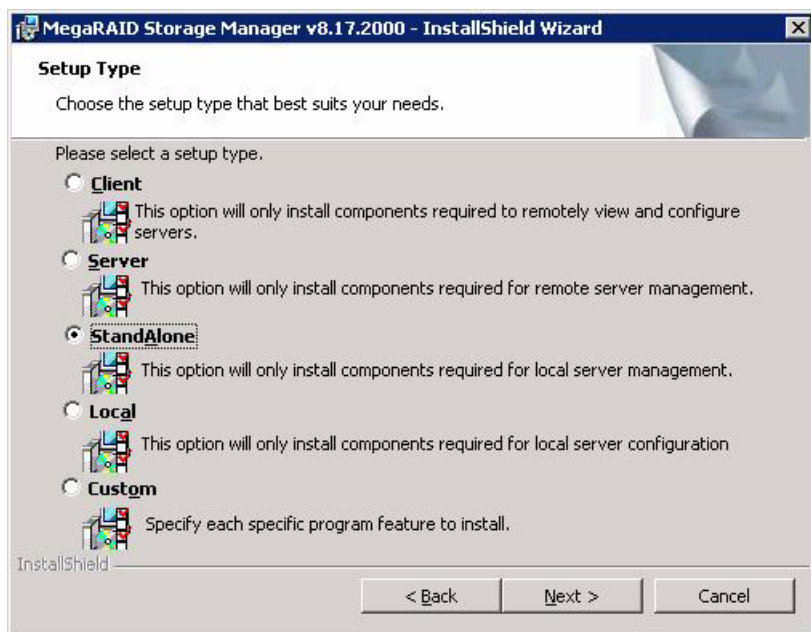


Figure 168: Setup Type Window

10. Select one of the setup options.
 - Select the **Client** radio button if you are installing the MegaRAID Storage Manager software on a computer that will be used to view and configure servers over a network. To begin installation, click **Install** on the next window that appears.

In the Client mode of installation, the MegaRAID Storage Manager software installs only client-related components, such as the MegaRAID Storage Manager GUI, and monitor configurator.

Use this mode when you want to manage and monitor servers remotely. When you install the MegaRAID Storage Manager software in Client mode on a laptop or a desktop, you can log in to a specific server by providing the IP address.
 - Select the **Server** radio button to install only those components required for remote server management. To begin installation, click on **Install** on the next window that appears.

- Select the **StandAlone** radio button if you will use the MegaRAID Storage Manager software to create and manage storage configurations on a standalone workstation. To begin installation, click on **Install** on the next window that appears.
- Select the **Local** radio button if you want to view only the workstation that has the MegaRAID Storage Manager software installed. You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.
- Select the **Custom** radio button if you want to specify individual program features to install.

If you select **Custom**, a window listing the installation features appears, as shown in [Figure 169](#). Select the features you want on this window.

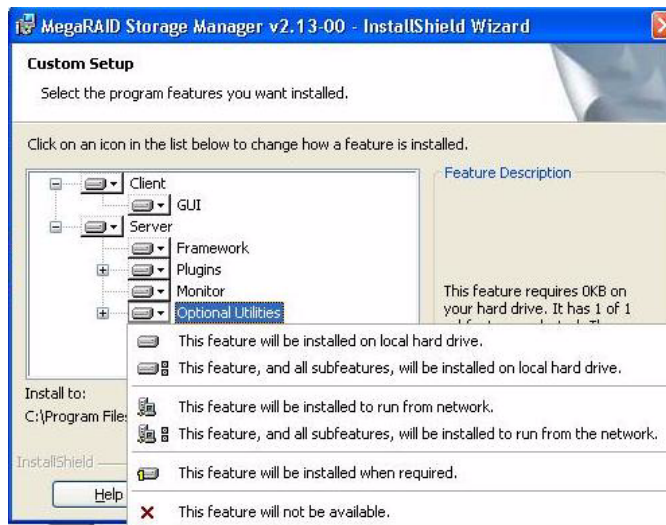


Figure 169: Custom Setup Window

11. Click **Next** to proceed.
12. Click **Install** to install the program.
13. When the final Configuration Wizard window appears, click **Finish**.

If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the server window appears. The MegaRAID Storage Manager - Host View window does not list any servers. You can use the MegaRAID Storage Manager - Host View window to manage systems remotely.

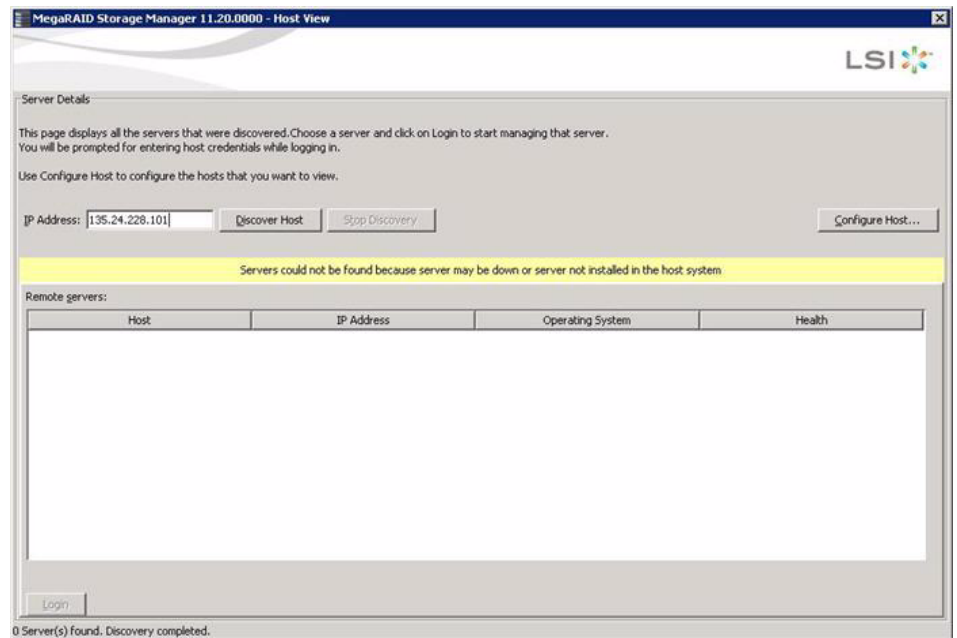


Figure 170: MegaRAID Storage Manager - Host View Window

6.4.3 Installing MegaRAID Storage Manager for the Solaris SPARC Operating System

Perform the following steps to install the MegaRAID storage Manager Software for Solaris 10 SPARC.

1. Copy the `MegaRaidStorageManager-SOLSPARC-8.10-.....tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOLSPARC-8.10-.....tar.gz` file using the following command:

```
tar -zxvf  
"MegaRaidStorageManager-SOLSPARC-8.10-.....tar.gz"
```

This step creates a new disk directory. Go to the new disk directory, and find and read the `readme.txt` file.

3. Enter the Bash shell.
4. Execute the command `./install.sh` present in the disk directory.
5. When prompted by the installation scripts, type `Y` to complete the installation.

NOTE: CacheCade Pro 2.0 SSD Read/Write Caching software is not applicable in SPARC.

6.4.4 Uninstalling MegaRAID Storage Manager Software for Solaris SPARC Operating System

To uninstall the MegaRAID Storage Manager software on a system running Solaris 10 SPARC, perform the following steps:

1. Execute the `Uninstaller.sh` file located in `/opt/MegaRaidStorageManager` directory.
2. When prompted by the uninstallation scripts, select `Y` to complete the installation.

NOTE: To shut down the MegaRAID Storage Manager Framework service, run the `svcadm disable -t MSMFramework`.

It is advisable to stop the Monitor service before stopping the MegaRAID Storage Manager Framework service.

To stop the Monitor service, run `svcadm disable -t MSMMonitor`.

To start the Framework service, run `svcadm enable MSMFramework`.

When the service is in maintenance state, run the `svcadm clear MSMFramework` command and `svcadm clear MSMMonitor` command.

To start the monitor service, run `svcadm enable MSMMonitor` command.

To check the status of the MegaRAID Storage Manager services, run the command `svcs -a | grep -i msm`.

6.4.5 Installing MegaRAID Storage Manager for Linux Operating System

Follow these steps if you need to install the MegaRAID Storage Manager software on a system running Red Hat Linux 3.0, 4.0, 5.0, 6.0 or SUSE Linux/SLES 9, 10, and 11:

1. Copy the `MSM_linux_installer-11.02.00-00.tar.gz` file to a temporary folder.
2. Untar the `MSM_linux_installer-11.02.00-00.tar.gz` file using the following command:

```
tar -zxvf MSM_linux_installer-11.02.00-00-...tar.gz
```

A new disk directory is created.

3. Go to the new `disk` directory.
4. In the `disk` directory, find and read the `readme.txt` file.
5. To start the installation, enter the following command:

```
cdsh install.sh -a
```

If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the MegaRAID Storage Manager - Host Name window appears. The MegaRAID Storage Manager - Host Name window does not list any servers. You can use this window to manage systems remotely.

To install the software using an interactive mode, execute the command `./install.sh` from the installation disk.

To install the product in a non-interactive or silent mode, use the command `./install.sh [-options] [-ru popup]` from the installation disk. The installation options are as follows:

- **Complete**
- **Client Component Only**
- **StandAlone**
- **Local**

The `-ru popup` command removes the pop-up from the installation list.

You also can run a non-interactive installation using the `RunRPM.sh` command.

The installer offers the following setup options:

- **Complete** – This installs all the features of the product.
- **Client Components Only** – The `storelib` feature of the product is not installed in this type of installation. As a result, the resident system can only administer and configure all of the servers in the subnet, but it cannot serve as a server.
- **StandAlone** – Only the networking feature will not be installed in this case. But the system can discover other servers in the subnet and can be discovered by the other servers in the subnet.
- **Local** - This option enables you to view only the workstation that has the MegaRAID Storage Manager software installed. You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.

This installation helps you select any of the setup types, but if you run `RunRPM.sh`, it installs the complete feature.

6.4.6 Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.0 x64 Operating System

Before installing the MegaRAID Storage Manager software on RHEL 6.0 x64 system, install the following RPMs. Without these RPMs the MegaRAID Storage Manager software might not install properly or might not work as expected.

- `libstdc++-4.4.4-13.el6.i686.rpm`
- `compat-libstdc++-33-3.2.3-69.i686.rpm`
- `libXau-1.0.5-1.el6.i686.rpm`
- `libxcb-1.5-1.el6.i686.rpm`
- `libX11-1.3-2.el6.i686.rpm`
- `libXext-1.1-3.el6.i686.rpm`
- `libXi-1.3-3.el6.i686.rpm`
- `libXtst-1.0.99.2-3.el6.i686.rpm`

The RHEL6.0 RC1 complete operating system installation is required for the MegaRAID Storage Manager software to work. The above mentioned rpm's come as part of RHEL6 RC1 Operating System DVD. These RPMs might need additional dependent RPMs as well, and you must also install all of the dependent RPMs on the target system.

6.4.7 Linux Error Messages

The following messages can appear while you are installing the MegaRAID Storage Manager software on a Linux operating system:

- More than one copy of MegaRAID Storage Manager software has been installed.

This message indicates that the user has installed more than one copy of the MegaRAID Storage Manager software. (This step can be done by using the `rpm-force` command to install the `rpm` file directly, which is not recommended, instead of using the `install.sh` file.) In such cases, the user must uninstall all of the `rpm` files manually before installing the MegaRAID Storage Manager software with the procedure listed previously.

- The version is already installed.

This message indicates that the version of the MegaRAID Storage Manager software you are trying to install is already installed on the system.

- The installed version is newer.

This message indicates that a version of the MegaRAID Storage Manager software is already installed on the system, and it is a newer version than the version you are trying to install.

- Exiting installation.

This is the message that appears when the installation is complete.

- RPM installation failed.

This message indicates that the installation failed for some reason. Additional message text explains the cause of the failure.

6.4.8 Kernel Upgrade

If you want to upgrade the kernel in the Linux operating system, you must restart the MegaRAID Storage Manager Framework and Services in the same order by entering the following commands.

```
/etc/init.d/vivaldiframeworkd restart
/etc/init.d/mrmonitor restart
```

6.4.9 Uninstalling MegaRAID Storage Manager Software on Linux

To uninstall the MegaRAID Storage Manager software on a system running Linux, follow these steps:

1. Go to `/usr/local/MegaRAID Storage Manager`.
2. Run `./uninstaller.sh`.

This procedure uninstalls the MegaRAID Storage Manager software.

6.4.9.1 Executing a CIM Plug-in on Red Hat Enterprise Linux 5

To execute a Common Information Model (CIM) plug-in on Red Hat Enterprise Linux 5, you must create the following symbolic links:

1. `cd /usr/lib` on RHEL 5
2. Search for `libcrypto`, `libssl`, and `libsysfs` libraries as follows:


```
ls -lrt libcrypto*, ls -lrt libssl*, ls -lrt libsysfs*
```
3. If the files `libcrypto.so.4`, `libssl.so.4`, and `libsysfs.so.1` are missing, manually create sym links as follows:


```
ln -s libcrypto.so libcrypto.so.4
ln -s libssl.so libssl.so.4
ln -s libsysfs.so libsysfs.so.1
```

For more information about CIM, see [MegaRAID Storage Manager Support on the VMware ESXi Operating System](#).

If the `.so` files are not present in the `/usr/lib` directory, create a link with the existing version of the library. For example, if `libcrypto.so.6` is present and `libcrypto.so` is not, create the link as follows:

```
ln -s libcrypto.so.6 libcrypto.so.4
```

6.4.10 MegaRAID Storage Manager Customization

On a 64-bit operating system, the system libraries are present in the `/usr/lib64` directory by default. However, for supporting CIM Plug-in, make sure that the libraries are also present in `/usr/lib` by installing the appropriate RPMs.

You can customize your Logo and Splash window by editing the `msm.properties` file present in the `<installation-directory>\MegaRAID Storage Manager` folder.

The `msm.properties` file has no values for the following keys:

- a. CHANNELLOGO=
- b. CHANNELSPLASHSCREEN=

No default values are assigned for these keys; therefore, the MegaRAID Storage Manager uses the default LSI Logo and splash screen.

To customize the Logo and splash screen, enter the Logo and Splash screen file name against these entries.

To enter the file names follow these steps:

1. Open the `msm.properties` file in the `<installation-directory>\MegaRAID Storage Manager` folder.
2. Enter the value for the logo file against the key CHANNELLOGO.
3. Enter the value for the splash screen file against the key CHANNELSPLASHSCREEN.
4. Save the file.
5. Place these two images in the `<installation-directory>\MegaRAID Storage Manager` folder.
6. Start the application.

Following are some of important points that you need to keep in mind:

- File names for both entries should not have any spaces. For example, the valid file name would be: `logo_test_1.png`, `LogoTest1.png`, or `TEST_SPLASH_FILE.jpg`.
- The logo image dimensions should not exceed 160 pixels x 85 pixels (width x height).
- The splash screen image dimensions should not exceed 390 pixels x 260 pixels (width x height).

After making the changes mentioned previously, when you log into the MegaRAID Storage Managers software, you will be able to view the changes with the new splash screen and logo in the MegaRAID Storage Manager software.

6.5 MegaRAID Storage Manager Support and Installation on VMware

6.5.1 Pre-requisites for Installing MegaRAID Storage Manager for VMware

This section documents the installation of the MegaRAID Storage Manager software on VMware ESX (also known as Classic) and on the VMware ESXi operating system.

For the VMware 3.5 operating system, it is necessary to install `libstdc++34-3.4.0-1.i386.rpm` before installing the MegaRAID Storage Manager software. You can download the rpm file from:

```
http://rpm.phone.net/index.php3/stat/4/idpl/1203252/com/1
libstdc++34-3.4.0-1.i386.rpm.html.
```

For the VMware 4.1 operating system, it is necessary to create a soft link as follows before installing the MegaRAID Storage Manager software. Run the following command to create the necessary soft link required for the MegaRAID Storage Manager software to work.

```
sudo ln -sf
/lib/libgcc_s.so.1/usr/lib/vmware/lib/libgcc_s.so.1
```

6.5.2 Installing MegaRAID Storage Manager on VMware ESX (VMware Classic)

The VMware operating system does not support any graphics components. To install the MegaRAID Storage Manager software on the VMware operating system, run the script `./vmware_install.sh` from the installation disk.

The installer lets you accept the license agreement, operating system, and storelib as follows:

- End user license agreement
- Operating system (VMware 3.5 or VMware 4.x operating system)
- Select the Storelib (Inbox Storelib or Storelib from the MegaRAID Storage Manager package)

NOTE: VMware Classic is not supported on VMware 5.x and higher versions.

6.5.3 Uninstalling MegaRAID Storage Manager for VMware

To uninstall the Server Component of the MegaRAID Storage Manager software on VMware, either use the `Uninstall` command in the Program menu, or run the script `/usr/local/MegaRAID Storage Manager/uninstaller.sh`.

You need to keep in mind the following points:

- A MegaRAID Storage Manager upgrade is supported in this release. Future releases can update this release.
- To shut down the MegaRAID Storage Manager Framework service, run the following command:

```
/etc/init.d/vivaldiframeworkd stop
```

It is recommended that you stop the Monitor service before you stop the MegaRAID Storage Manager Framework service. To stop the Monitor service run the following command:

```
/etc/init.d/mrmonitor stop
```

The Linux RPM of the MegaRAID Storage Manager software works under the console with minimal changes. Hardware RAID is currently supported in ESX 4.x.

NOTE: There is a known limitation that virtual drives that are created or deleted will not be reflected to the kernel. The workaround is to reboot the server or to run `esxcfg-rescan <vmhba#>` from COS shell.

6.5.4 MegaRAID Storage Manager Support on the VMware ESXi Operating System

This section outlines the product requirements needed to support the VMware ESXi operating system. Classic VMware includes a service console that is derived from the Linux 2.4 kernel, but with reduced functionality.

The MegaRAID Storage Manager server part cannot be installed directly in the VMware ESXi operating system. Management is possible only through Common Information Model (CIM) providers. These CIM providers integrated into the VMware ESXi system build an interface between the hardware driver of the LSI MegaRAID controller and remote applications, such as the MegaRAID Storage Manager software. Management is performed through the MegaRAID Storage Manager software installed on a remote machine (Linux/Windows).

VMware ESXi comes with the Small Footprint CIM Broker (CFCB) CIM Object Manager (or CIMOM). A CIMOM manages the communication between providers, which interacts with the hardware, and a CIM client, where the administrator manages the system.

SFCB supports Common Manageability Programming Interface (CMPI)-style providers. CMPI defines a common standard used to interface Manageability Instrumentation (providers, instrumentation) to Management Brokers (CIM Object Manager). CMPI standardizes Manageability Instrumentation, which lets you write and build instrumentation once and run it in different CIM environments (on one platform).

Network communication is a key element for a proper communication between the ESXi CIM provider and the LSI management software. Please make sure that the network settings are correct by making the following changes:

- For networks that do not have DNS configured, the “hosts” file in the machine on which the MegaRAID Storage Manager software is installed must be edited as follows:
 - a. Add an entry to map the VMware host’s IP address with the host name. This is for the discovery process to happen correctly. In the absence of this entry, the VMware host would be discovered as 0.0.0.0.
 - b. Add an entry to map its own IP address (not the loop back address) with the Hostname. This is to ensure that the Asynchronous Event Notifications (AENs) are delivered correctly.
- For networks that have DNS configured, perform the following:
 - a. Provide the correct DNS server IP address when you do the initial configurations for the VMware host.
 - b. In the hosts file of the machine on which the MegaRAID Storage Manager software is installed, add an entry to map its own IP address (not the loop back address) with the host name, to ensure that the Asynchronous Event Notifications are delivered correctly.

6.5.5 Limitations

The following are the limitations of this installation and configuration.

- No status information exists for the controller

6.5.5.1 Differences in MegaRAID Storage Manager for the VMware ESXi System

- Events are collected as long as the MegaRAID Storage Manager software runs on the client.
- The MegaRAID Storage Manager software on VMware responds slower as compared to the response of the MegaRAID Storage Manager software on Windows/Linux/Solaris. Events are collected from the time a client logs in to an ESXi machine for the first time, and it continues to be collected as long as the Framework is running.

The following are some of the differences in the MegaRAID Storage Manager utility when you manage a VMware server.

- The following limitations apply to the system information exposed through the application:
 - Only the IP address and the host name appear.
 - The operating system type and the operating system architecture do not appear.
 - No support exists for the controller health information.
- Authentication support:
 - The MegaRAID Storage Manager software allows CIMOM server authentication with the user ID and the password for VMware.
 - Access to VMware ESXi hosts is controlled based on the user privileges. Only root users can have full access, while the non-root users can have only view only access.
 - Multiple root users can simultaneously login using 'Full Access' mode to access the VMware ESXi server.

- Event logging:

The event logging feature for the MegaRAID Storage Manager Client connected to a VMware ESXi system behaves as follows:

- No support exists for retrieving initial logs (the events that occurred before a client logs in for the first time). Only those events that occur after the client logs in appear in the Event Logger dialog.
- The ESXi system log is displayed in the system in which the MegaRAID Software Manager Framework runs.
- The View Log option allows you to view the logs saved in a text file on the Event Logger dialog.
- Refreshing of the MegaRAID Storage Manager GUI after any updates on the firmware is slower for a client connected to VMware ESXi hosts, compared to one that is connected to the Windows, Linux, or Solaris hosts.
- It takes time to discover the CIMOM servers. If you start the MegaRAID Storage Manager client immediately after you install the MegaRAID Storage Manager software (or restart Framework service), you will not be able to discover any hosts in the network.
- VMware ESXi is supported only on a full installation of the MegaRAID Storage Manager software; standalone, client-only, server-only, and local modes do not support VMware ESXi management.
- VMware ESXi is supported on following operating systems:
 - Microsoft Windows Server

- RHEL
- SuSE Linux

6.6 Installing and Configuring a CIM Provider

This section describes the installation and configuration of the LSI MegaRAID Common Information Model (CIM) provider. The Common Information Model offers common definitions of management information for networks, applications, and services, and allows you to exchange management information across systems throughout a network.

On a VMware ESXi system, management is possible only through a CIM provider, and it is performed through the MegaRAID Storage Manager software installed on a remote machine running a Linux or Windows operating system.

The VMware ESXi system comes with the Small Footprint CIM Broker (SFCB) CIM Object Manager (or CIMOM). A CIMOM manages communication between providers, which interact with the hardware, and a CIM client, where the administrator manages the system.

SFCB supports Common Manageability Programming Interface (CMPI)-style providers. CMPI defines a common standard used to interface manageability instrumentation (providers, instrumentation) to management brokers (CIM Object Manager). CMPI standardizes manageability instrumentation, which allows you to write and build instrumentation once and run it in different CIM environments (on one platform).

6.6.1 Installing a CIM SAS Storage Provider on the Linux Operating System

The following procedure documents how to install and uninstall the LSI CIM SAS Storage Provider on a system running on the Linux operating system.

NOTE: Uninstall all the previous versions of LSI SAS Provider before you install this version. You can check all of the installed versions of LSI SAS Provider by running the `rpm -qa | grep LsiSASProvider` command.

- To install a CIM SAS Storage Provider on a Linux system, install the SAS Provider using the Red Hat Package Manager (RPM) by entering the following command:

```
rpm -ivh
```

The RPM installs all of the necessary files and the Managed Object Format (MOF), and it registers the libraries. The SAS Provider is now ready to use.

NOTE: After you install LSI CIM SAS Provider, the MOF file `LSI_SASRaid.mof` is available under the `/etc/lsi_cimprov/sas/pegasus/common` directory.

- To uninstall a CIM SAS Storage Provider on a Linux system, remove LSI CIM SAS Provider by entering the command:

```
rpm -ivh LsiSASProvider-<version>.<arch>.rpm
```

This removes all of the necessary files, uninstalls the MOF, and unregisters the libraries. The SAS Provider is no longer on the system.

NOTE: Tog-pegasus binaries, such as cimmoof, cimprovider, and wbemexec, should be in the PATH variable of `/etc/profile`, and hence, are defined in all environments of the system.

To run the CIM SAS Storage Provider on Pegasus version 2.5.x, perform the following steps:

1. After you install the LSI SAS Pegasus provider, verify that the `libLsiSASProvider.so` file and the `libLsiSASProvider.so.1` file are in `/usr/lib/Pegasus/providers` directory.

If these files are not present, copy the `libLsiSASProvider.so.1` file from `/opt/tog-pegasus/providers/lib` to `/usr/lib/Pegasus/providers`, and create a symbolic link `libLsiSASProvider.so` to `/usr/lib/Pegasus/providers/libLsiSASProvider.so.1` at `/usr/bin/Pegasus/providers`.

2. Restart the Pegasus CIM Server and LSIServer by performing the following steps:
 - To start the tog-pegasus server, run the following command:


```
# /etc/init.d/tog-pegasus restart
```
 - To start LSISAS Sever, run the following command:


```
# /etc/init.d/LsiSASd restart
```

6.6.2 Installing a CIM SAS Storage Provider on Windows

The following procedure describes how to install and uninstall the LSI CIM SAS Storage Provider on a system running on a Windows operating system.

Perform the following steps to install a CIM SAS Storage Provider on a Windows system:

1. Go to DISK1.
2. Run `setup.exe`.

The installer installs all of the necessary files and the MOF, and registers the COM DLL. The CIM SAS Provider is now ready to use.

Perform the following steps to uninstall a CIM SAS Storage Provider on a Windows operating system.

1. Select **Control Panel > Add/Remove Program**.
2. Remove the LSI WMI SAS Provider Package.

This step removes all of the necessary files, uninstalls the MOF, and unregisters the COM dll. The SAS Provider is no longer on the system.

6.7 Installing and Configuring an SNMP Agent

A Simple Network Management Protocol (SNMP)-based management application can monitor and manage devices through SNMP extension agents. The MegaRAID SNMP subagent reports the information about the RAID controller, virtual drives, physical devices, enclosures, and other items per SNMP request. The SNMP application monitors these devices for issues that might require administrative attention.

This section describes the installation and configuration of the LSI MegaRAID SNMP agent on Linux, Solaris, and Windows operating systems.

NOTE: The complete installation of the MegaRAID Storage Manager software installs the SNMP agent. However, you can install the SNMP agent (installer) on a system separately, without the MegaRAID Storage Manager software being installed

6.7.1 Prerequisite for LSI SNMP Agent RPM Installation

The LSI SNMP agent application depends upon the standard SNMP Utils package. Make sure that the SNMP-Util package is present in the system before you install LSI SNMP agent RPM.

The SNMP-Util package includes the RPM's net-snmp-libs, net-snmp-utils, and additional dependent RPMs.

Make sure that these RPM's are installed from the operating system media before you install the LSI SNMP agent RPM.

6.7.2 Prerequisite for Installing SNMP Agent on Linux Server

The SNMP application requires the standard library libsysfs. Make sure that this library is present in the system before installing the SNMP RPM.

The minimum library versions required for installing SNMP server are as follows.

- libsysfs version 2.0. This library is available in the rpm `<Lib_Utils-1.xx-xx.noarch.rpm>`.
`<Lib_Utils-1.xx-xx.noarch.rpm>` is packaged in the SNMP zip file.
- `libstdc++.so.6`. This library is present in `/usr/lib` directory.

You can install the SNMP application from the Linux software component RPM that provides these libraries. These RPM's are available in the Linux OS DVD.

6.7.3 Installing and Configuring an SNMP Agent on a Linux Operating System

This section explains how to install and configure the SAS SNMP Agent for the SUSE Linux and Red Hat Linux operating systems.

Perform the following steps to install and configure the SAS SNMP Agent for the SUSE Linux and Red Hat Linux operating systems:

NOTE: This procedure requires that you have the Net-SNMP agent installed on the Linux machine.

NOTE: The RPM has not been created to support -U version. The RPM -U will probably fail with this RPM.

1. Install the LSI SAS SNMP Agent using the `rpm -ivh <sas rpm>` command.

NOTE: After installation, find the SAS MIB file `LSI-AdapterSAS.mib` under the `/etc/lsi_mrdsnm/sas` directory.

RPM makes the necessary modification needed in the `snmpd.conf` file to run the agent.

NOTE: Before installation, check whether there is any pass command exists that starts with 1.3.6.1.4.1.3582 OID in `snmpd.conf`. If so, delete all of the old pass commands that start with 1.3.6.1.4.1.3582 OID. (This situation could occur if an earlier version of LSI SNMP Agent was installed in the system.)

The `snmpd.conf` file structure should be the same as the file structure `lsi_mrdsnmpd.conf`. For reference, a sample configuration file (`lsi_mrdsnmpd.conf`) is in the `/etc/lsi_mrdsnmp` directory.

2. To run an SNMP query from a remote machine, add the IP address of that machine in the `snmpd.conf` file, as in this example:

```
com2sec      snmpclient      172.28.136.112      public
```

Here, the IP address of the remote machine is 172.28.136.112.

3. To receive an SNMP trap to a particular machine, add the IP address of that machine in the `com2sec` section of the `snmpd.conf` file.

For example, to get a trap in 10.0.0.144, add the following to `snmpd.conf`.

```
#          sec.name      source      community
com2sec    snmpclient      10.0.0.144      public
```

4. To run or stop the `snmpd` daemon, enter the following command:

```
/etc/init.d/snmpd start
/etc/init.d/snmpd stop
```

5. To start/stop the SAS SNMP Agent daemon before issuing a SNMP query, enter the following command:

```
/etc/init.d/lsi_mrdsnmpd start
/etc/init.d/lsi_mrdsnmpd stop
```

You can check the status of the SAS SNMP Agent daemon by checked by entering the following command:

```
/etc/init.d/lsi_mrdsnmpd status
```

6. Issue an SNMP query in this format:

```
snmpwalk -v1 -c public localhost .1.3.6.1.4.1.3582
```

7. You can get the SNMP trap from local machine by issuing the following command:

```
snmptrapd -P -F "%02.2h:%02.2j TRAP%w.%q from %A %v\n"
```

NOTE: To receive a trap in a local machine with Net-SNMP version 5.3, you must modify the `snmptrapd.conf`, file (generally located at `/var/net-snmp/snmptrapd.conf`). Add `disableAuthorization yes` in `snmptrapd.conf` and then run `sudo snmptrapd -P -F "%02.2h:%02.2j TRAP%w.%q from %A %v\n"`.

NOTE: It is assumed that `snmpd.conf` is located in `/etc/snmp` for the Red Hat operating system and `/etc` for the SLES operating system. You can change the file location from the `/etc/init.d/lsi_mrdsnmpd` file.

You can install SNMP without the trap functionality. To do so, set the `TRAPIND` environment variable to "N" before running RPM.

Before you install a new version, you must uninstall all previous versions.

For the SLES 10 operating system, perform the following steps to run SNMP:

1. Copy `/etc/snmp/snmpd.conf` to `/etc/snmpd.conf`.
2. Modify the `/etc/init.d/snmpd` file, and change `SNMPDCONF=/etc/snmp/snmpd.conf` entry to `SNMPDCONF=/etc/snmpd.conf`.
3. Run `LSI SNMP rpm`.

6.7.4 Installing and Configuring an SNMP Agent on the Solaris Operating System

6.7.4.1 Prerequisites

This section explains how to install and configure SAS SNMP Agent for the Solaris operating system.

This package requires that you have Solaris System Management Agent installed on the Solaris machine.

6.7.4.2 Installation SNMP on the Solaris Operating System

To install SNMP for the Solaris operating system, perform the following steps:

1. Unzip the LSI SAS SNMP Agent package.
2. Run the install script by using the following command:

```
# ./install.sh
```

NOTE: The installation exits if any existing versions of `storelib` and `sassnmp` are installed on the Solaris machine. Uninstall the existing version by using the following commands:

```
# pkgrm storelib (to uninstall the storelib library)
# pkgrm sassnmp (to uninstall the LSI SAS SNMP Agent)
```

6.7.4.3 LSI SAS SNMP MIB Location

After you install the LSI SAS SNMP Agent package, the MIB file `LSI-AdapterSAS.mib` is installed under `/etc/lsi_mrdsnmp/sas` directory.

6.7.4.4 Starting, Stopping, and Checking the Status of the LSI SAS SNMP Agent

The following commands are used to start, stop, restart, and check the status of the Solaris System Management Agent (`net snmpd`) daemon:

- Start: # `svcadm enable svc:/application/management/sma:default`
- Stop: # `svcadm disable svc:/application/management/sma:default`
- Restart: # `svcadm restart svc:/application/management/sma:default`
- Status: # `svcs svc:/application/management/sma:default`

NOTE: `Online` indicates that the SMA is started. `Disabled` indicates that the SMA is stopped.

The following commands are used to start, stop, restart, and check the status of the SAS SNMP Agent daemon:

- Start: `#!/etc/init.d/lsi_mrdsnmpd start`
- Stop: `#!/etc/init.d/lsi_mrdsnmpd stop`
- Restart: `#!/etc/init.d/lsi_mrdsnmpd restart`
- Status: `#!/etc/init.d/lsi_mrdsnmpd status`

6.7.4.5 Configuring the `snmpd.conf` File

By default, you can run the SNMP queries (walk, get) from any remote machine without any changes to the `snmpd.conf` file. To quickly add a new community and client access, perform the following steps:

1. Stop the SMA service by running the following command:

```
# svcadm disable svc:/application/management/sma:default
```

2. Add read-only and read-write community names.

- a. Add a read-only community name and client/hostname/ipaddress under SECTION: Access Control Setup in the `/etc/sma/snmp/snmpd.conf` file, as shown in the following excerpt.

```
#####
# SECTION: Access Control Setup
# This section defines who is allowed to talk to
# your running SNMP Agent.
# rocommunity: a SNMPv1/SNMPv2c read-only access
# community name
# arguments: community
# [default|hostname|network/bits] [oid]
# rocommunity snmpclient 172.28.157.149
#####
```

- b. Add a readwrite community name and client, hostname, ipaddress under SECTION: Access Control Setup in `/etc/sma/snmp/snmpd.conf` file, as shown in the following excerpt.

```
#####
# SECTION: Access Control Setup
# This section defines who is allowed to talk to your
# running snmp agent.
# rocommunity: a SNMPv1/SNMPv2c read-only access
# community name
# arguments: community
# [default|hostname|network/bits] [oid]
# rwcommunity snmpclient 172.28.157.149
#####
```

3. Start the SMA service by using the following command:

```
# svcadm enable svc:/application/management/sma:default
```

NOTE: Refer to the command `man snmpd.conf` for more information about configuring the `snmpd.conf` file.

6.7.4.6 Configuring SNMP Traps

To receive SNMP traps, perform the following steps:

1. Stop the LSI SAS SNMP Agent by using the following command:

```
#/etc/init.d/lsi_mrdsnmpd stop
```

2. Edit the `/etc/lsi_mrdsnmp/sas/sas_TrapDestination.conf` file, and add the Ip address as shown in the following excerpt.

```
#####
# Agent Service needs the IP addresses to sent trap
# The trap destination may be specified in this file
# or using snmpd.conf file. Following indicators can
# be set on "TrapDestInd" to instruct the agent to
# pick the IPs as the destination.
# 1 - IPs only from snmpd.conf
# 2 - IPs from this file only
# 3 - IPs from both the files
#####
TrapDestInd 2
##### Trap Destination IP #####
127.0.0.1    public
172.28.157.149 public
#####
```

3. Start the LSI SAS SNMP Agent by entering the following command:

```
#/etc/init.d/lsi_mrdsnmpd start
```

6.7.4.7 Uninstalling the SNMP Package

The `uninstall.sh` script is located under the `/etc/lsi_mrdsnmp/sas` directory. Use the following command to uninstall the package:

```
# cd /etc/lsi_mrdsnmp/sas
# ./uninstall.sh
```

6.7.5 Installing an SNMP Agent on the Windows Operating System

This section explains how to install and configure SAS SNMP Agent for the Windows operating system.

6.7.5.1 Installing SNMP Agent

Perform the following steps to install SNMP Agent:

1. Run `setup.exe` from DISK1.
2. Use SNMP Manager to retrieve the SAS data (it is assumed that you have compiled LSI-AdapterSAS.mib file already).

The LSI-AdapterSAS.mib file is available under the `%ProgramFiles%\LSI Corporation\SNMPAgent\SAS` directory.

3. Use a trap utility to get the traps.

NOTE: Before you install the Agent, make sure that SNMP Service is already installed in the system.

6.7.5.2 Installing SNMP Service for the Windows Operating System

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for a Windows system.

1. Select **Add/Remove Programs** from the Control Panel.

2. Select **Add/Remove Windows Components** in the left side of the Add/Remove Programs window.
3. Select **Management and Monitoring Tools**.
4. Click **Next**, and follow any prompts to complete the installation procedure.

6.7.5.3 Configuring SNMP Service on the Server Side

Perform the following steps to configure SNMP Service on the server side.

1. Select **Administrative Tools** from the Control Panel.
2. Select **Services** in the Administrative Tools window.
3. Select **SNMP Service** in the Services window.
4. Open **SNMP Service**.
5. Click the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

6.7.5.4 Installing SNMP Service for the Windows 2008 Operating System

Before you install the LSI Agent, make sure that SNMP Service is already installed in the system.

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for Windows 2008 system.

1. Select **Program and Features** from the Control Panel.
2. Click **Turn windows feature on/off** to select the windows components to install.
3. Select **Features** from the menu.
4. Click **Add Features**.
5. Select **SNMP Services**.
6. Click **Next**.
7. Click **Install**, and the SNMP installation starts. You will be prompted for the Windows 2008 CD during the installation.
8. Insert the CD, and click **Ok**. The installation resumes.

After the installation is finished, the system displays a message saying that the installation is successful.

6.7.5.5 Configuring SNMP Service on the Server Side for the Windows 2008 Operating System

To configure SNMP service on the server side for Windows 2008 operating system, perform the following steps:

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** from Administrative Tools window.
3. Select **SNMP Service** from the Services window.
4. Open **SNMP Service**, and go to its properties.
5. Go to the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

6.8 MegaRAID Storage Manager Support and Installation on the Solaris 10 Operating System

6.8.1 Installing MegaRAID Storage Manager Software for the Solaris 10 Operating System

This section documents the installation of the MegaRAID Storage Manager software on the Solaris 10 U5, U6, U7, U8 x86 and x64 operating systems.

Follow these steps to install the MegaRAID Storage Manager software on a system running the Solaris 10 x86 operating system:

1. Copy the `MegaRaidStorageManager-SOLX86-....tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOLX86-....tar.gz` file using the following command:

```
tar -zxvf MegaRaidStorageManager-SOLX86-....tar.gz
```

This step creates a new disk directory.

3. Go to the new disk directory, and find and read the `readme.txt` file.
4. Enter the Bash shell.
5. Execute the command `./install.sh` present in the disk directory.
6. When prompted by the installation scripts, select `Y` to complete the installation.

6.8.2 Uninstalling MegaRAID Storage Manager Software for the Solaris 10 Operating System

Follow these steps to uninstall the MegaRAID Storage Manager software on a system running the Solaris 10 x86 operating system:

1. Execute the `Uninstaller.sh` file located in `/opt/MegaRaidStorageManager` directory.
2. When prompted by the uninstallation scripts, select `Y` to complete the installation.

NOTE: To shut down the MegaRAID Storage Manager Framework service, run `svcadm disable -t MSMFramework`. It is advisable to stop the Monitor service before stopping the MegaRAID Storage Manager Framework service. To stop Monitor service, run `svcadm disable -t MSMMonitor`.

NOTE: To start the Framework service, run `svcadm enable MSMFramework`. To start the monitor service, run `svcadm enable MSMMonitor`.

NOTE: When the service is in maintenance state, run `svcadm clear MSMFramework` and `svcadm clear MSMMonitor`.

NOTE: To check the status of MegaRAID Storage Manager services run the command `svcs -a | grep -i msm`.

6.9 Installing MegaCLI for VMware 5.0

MegaCLI is packaged into a vSphere Installation Bundle (VIB) for VMware 5.0.

To install the VIB, use the following command:

```
esxcli software vib install -v=<path to CLI VIB> --force
--maintenance-mode --no-sig-check
```

In the above command, the parameters specify the following information:

- `force` - Bypasses checks for package dependencies, conflicts, obsolescence, and acceptance levels.
- `maintenance-mode` - Pretends that the maintenance mode is in effect. Otherwise, the installation will stop for live installs that require maintenance mode.
- `no-sig-check` - Bypasses acceptance level verification, including signing.

For example, if the MegaCLI VIB, `vmware-esx-MegaCli-8.02.14.vib`, is present in the `/tmp` directory, it can be installed using the following command:

```
esxcli software vib install -v
=/tmp/vmware-esx-MegaCli-8.02.14.vib --force
--maintenance-mode --no-sig-check
```

Once the MegaCLI package is installed, it will be available in the `/opt/lsi/MegaCLI` directory.

To uninstall the VIB, use the following command:

```
esxcli software vib remove --force -n=<name of the VIB>
```

For example, if you are uninstalling the `vmware-esx-MegaCli-8.02.14.vib`, use the following command:

```
esxcli software vib remove --force
-n=vmware-esx-MegaCli-8.02.14
```

6.10 MegaRAID Storage Manager Remotely Connecting to VMware ESX

When the MegaRAID Storage Manager software is used to connect to a VMware ESX machine from a remote machine (Windows/Linux), we recommend that for long running operations (such as volume creation, deletion) to complete in a shorter time, you perform the following steps:

1. Login to the VMware ESX machine.
2. Open `/etc/sfcb/sfcb.cfg`.
3. Increase the `keepaliveTimeout` value from 1 to 100 or to a higher value.
4. Restart `sfcbd` (`/etc/init.d/sfcbd-watchdog restart`).
5. Restart the MegaRAID Storage Manager Framework on the MegaRAID Storage Manager client machine.
 - For Windows - restart the framework service.
 - For Linux - restart the `vivaldi` framework service.
6. Relaunch the MegaRAID Storage Manager window.

6.11 Prerequisites to Running MegaRAID Storage Manager Remote Administration

The MegaRAID Storage Manager software requires ports 3071 and 5571 to be open to function. Follow these steps to prepare to run the MegaRAID Storage Manager Remote Administration.

1. Configure the system with a valid IP address.

Make sure the IP address does not conflict with another in the sub network.

Ports, such as 3071 and 5571, are open and available for the MegaRAID Storage Manager framework communication.

2. Disable all security manager and firewall.
3. Configure the multicasting.

Make sure Class D multicast IP addresses are registered (at least 229.111.112.12 should be registered for the MegaRAID Storage Manager software to work); if not, create a static route using the following command:

```
Route add 229.111.112.12 dev eth1
```

4. Install the MegaRAID Storage Manager software. If the MegaRAID Storage Manager software is already installed, restart the MegaRAID Storage Manager Framework.

Chapter 7

MegaRAID Storage Manager Window and Menus

This chapter explains how to start the MegaRAID Storage Manager software and describes the MegaRAID Storage Manager window and menus.

7.1 Starting the MegaRAID Storage Manager Software

You must have administrative privileges to use the MegaRAID® Storage Manager software in either full-access or in view-only mode. Follow these steps to start the MegaRAID Storage Manager software on various platforms.

1. Start the program using the method required for your operating system environment:
 - To start the MegaRAID Storage Manager software on a Microsoft Windows operating system, select **Start >> Programs >> MegaRAID Storage Manager >> StartupUI**, or double-click the MegaRAID Storage Manager shortcut on the desktop.

NOTE: If a warning appears stating that Windows firewall has blocked some features of the program, click **Unblock** to allow the MegaRAID Storage Manager software to start. (The Windows firewall sometimes blocks the operation of programs that use JavaTechnology.)

- To start the MegaRAID Storage Manager software on a Red Hat Linux operating system, select **Applications>>System Tools >> MegaRAID Storage Manager StartupUI**.
- To start MegaRAID Storage Manager software on a SUSE Linux or SLES operating system, select **Start >> System >> More Programs >> MegaRAID Storage Manager**.
- To start MegaRAID Storage Manager software on a Solaris X86 and Solaris SPARC operating system, select **Launch >> Applications >> Utilities >> MegaRAID Storage Manager StartupUI**

When the program starts, the Select Server window appears, as shown in [Figure 171](#). The remote servers display, along with their IP addresses, operating system, and health status.

NOTE: If you do a local mode installation, as shown in [Section 6.4.2, Installing MegaRAID Storage Manager Software on Microsoft Windows, on page 281](#), [Figure 171](#) will not be displayed. It will directly prompt the user to the login dialog as shown in [Figure 173](#).

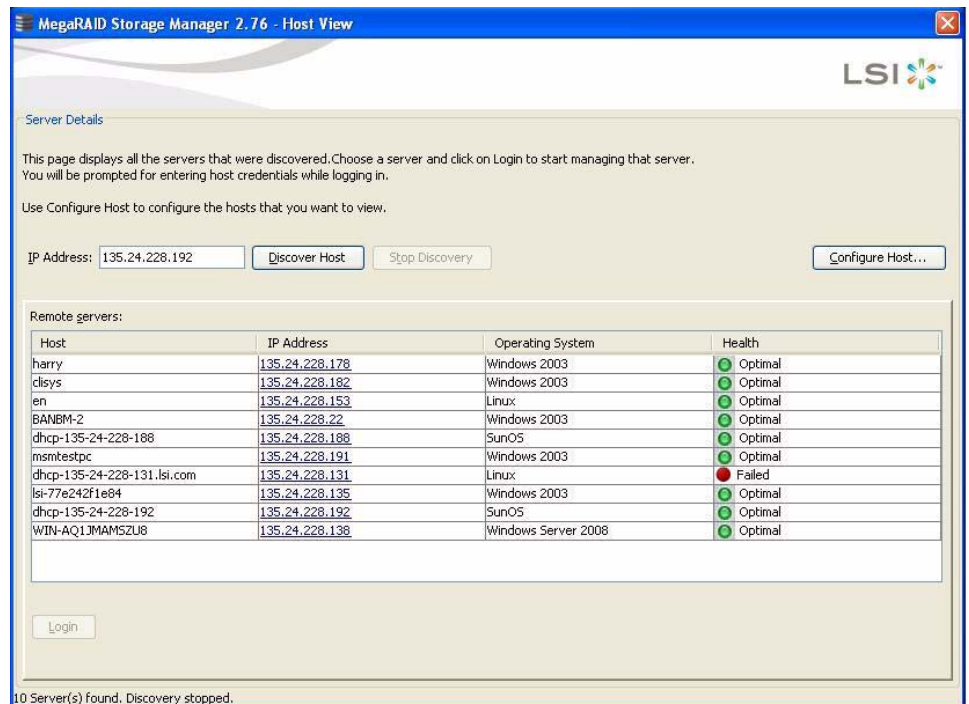


Figure 171: Select Server Window

The Select Server window shows an icon for each server on which the MegaRAID Storage Manager software is installed. The servers are color-coded with the following definitions:

- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (possibly because a drive in a virtual drive has failed).
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

NOTE: Do not enter the VMware ESXi server's IP address in the **IP Address** field in [Figure 171](#). Instead enter a valid MegaRAID Storage Manager server's IP address and select the **Display all the systems in the Network of the local server** option in [Figure 172](#).

If you click **Configure Host**, the configure host opens, as shown in the following figure.

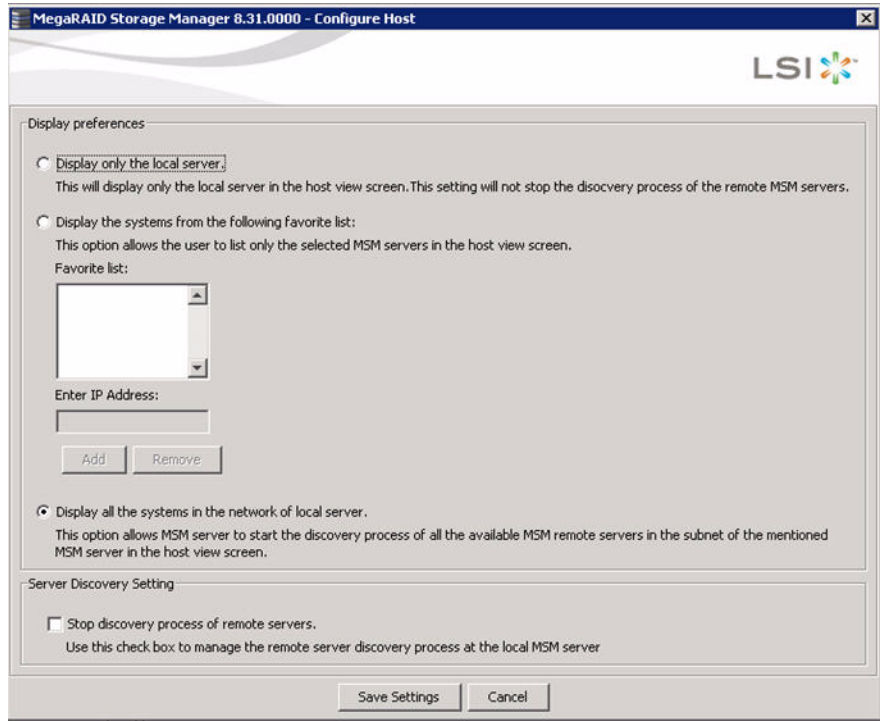


Figure 172: Configure Host

The following options are available to configure the host.

- Display only the local server – Select this option to display only the Local server or the Server of the IP address entered in the Host View screen.
- Display the systems from the following favorite list – You can enter the IP Addresses of the servers of choice to discover. It will discover only those servers.
- Display all the systems in the Network of the local server – Select this option to allow MegaRAID Storage Manager server to start the discovery process of all the available remote servers (including the VMware ESXi server) in the subnet mentioned MegaRAID Storage Manager server in the host view screen.

After selecting the configure host option click on **Save Settings** to save your setting, or on **Cancel** to quit without saving.

Select the **Stop discovery process of remote servers** check box and click on **Save Settings**, to abort the discovery process which has already begun. This check box is enabled only when there is a active discovery process.

The discovery process might take very long time to complete, you can use this function to abort this process.

NOTE: For the VMware ESXi, the server icon does not denote the health of the server. The icon is always green regardless of the health of the system.

NOTE: The VMware server does not show the system health and the operating system labels. It shows only the host name and the IP address of the server.

NOTE: When connecting to a VMware server on a different subnet, one or more frameworks have to be running in the subnet to connect to the CIMOM.

2. Double-click the icon of the server that you want to access.

The Server Login window appears, as shown in the following figure.



Figure 173: Server Login Window

3. Enter your user name and password.

The question mark icon opens a dialog box that explains what you need for full access to the server and for view-only access to the server. You will be allowed three attempts to Log in.

NOTE: When connected to VMware operating system, the Server Login screen shows only one label for access, Full Access. Multiple users can have full access to the VMware server.

4. Select an access mode from the drop-down menu for **Login Mode**, and click **Login**.
 - Select **Full Access** if you need to both view and change the current configuration.
 - Select **View Only** if you need to only view and monitor the current configuration.

NOTE: If the computer is networked, this login is for the computer itself, not the network login.

5. Enter the root or administrator user name and password to use Full Access mode.

NOTE: In Linux, users belonging to the root group can log in. You do not have to be the user root.

If your user name and password are correct for the Login mode you have chosen, the MegaRAID Storage Manager main menu appears.

7.2 MegaRAID Storage Manager Main Menu

This section describes the MegaRAID Storage Manager main menu, which is shown in Figure 174.

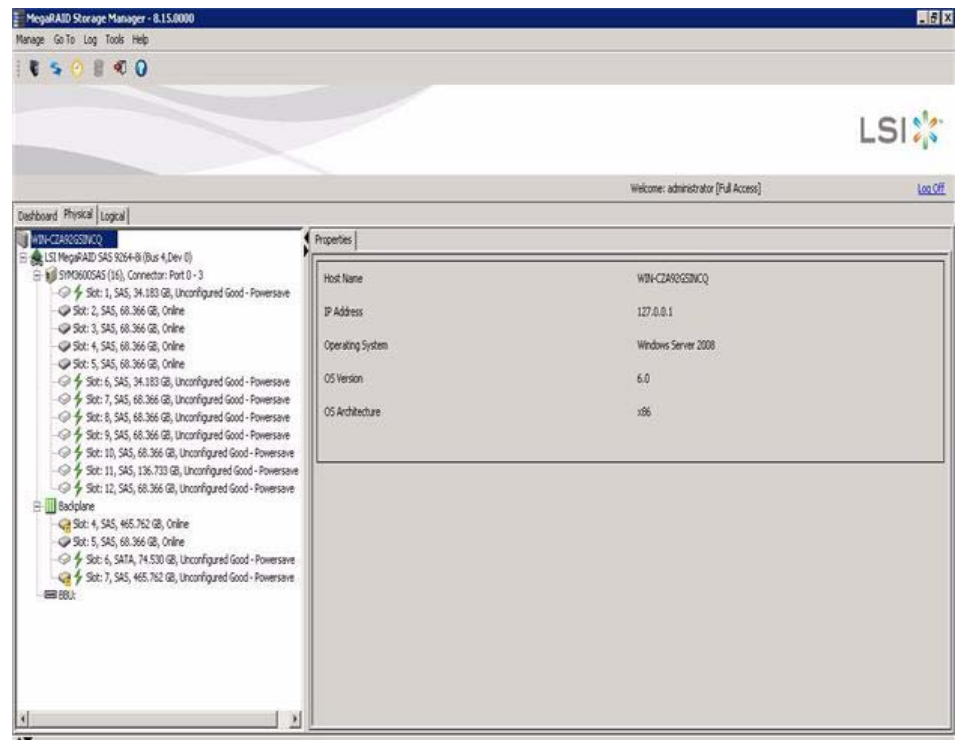


Figure 174: MegaRAID Storage Manager Main Window

The following topics describe the panels and menu options that appear on this screen.

7.2.1 Dashboard, Physical View, Logical View

7.2.1.1 Dashboard View

The left panel of the MegaRAID Storage Manager window displays the *Dashboard* view, the *Physical* view, or the *Logical* view of the system and the attached devices, depending on which tab is selected.

The *Dashboard* view shows an overview of the system and covers the following features:

- Properties of the virtual drives and the physical drives
- Total capacity, configured capacity, and unconfigured capacity
- Background operations in progress
- The MegaRAID Storage Manager software features and their status (enabled or disabled)
- Actions you can perform, such as creating a virtual drive and updating the firmware
- Links to online help

Figure 175 shows the Dashboard view.

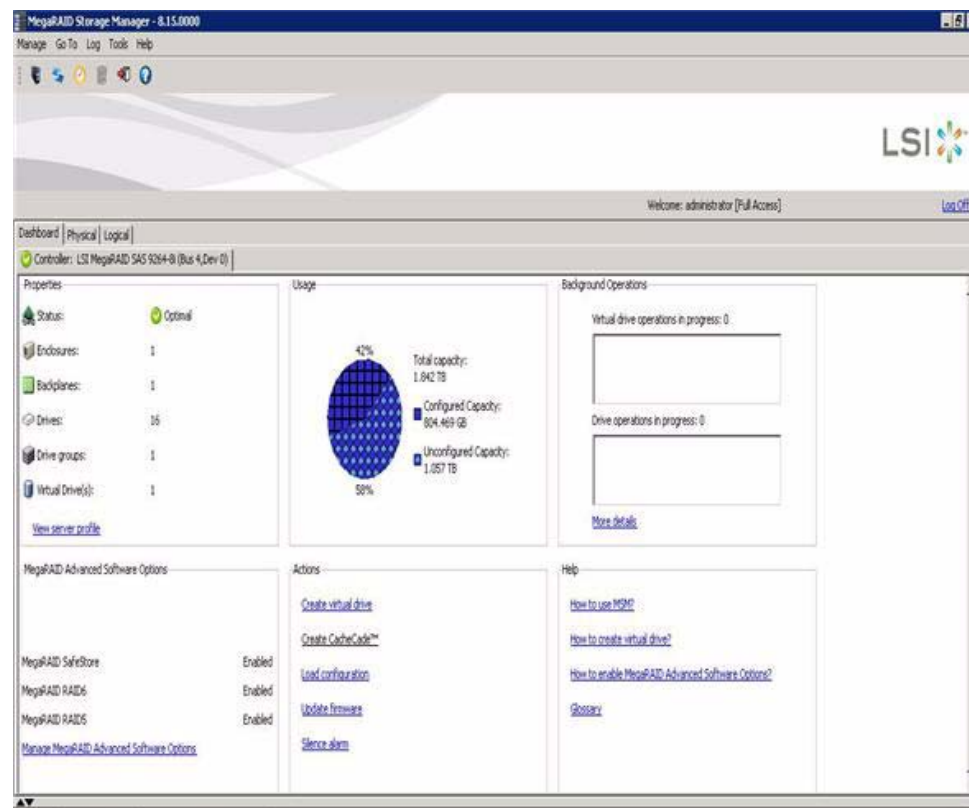


Figure 175: MegaRAID Storage Manager Dashboard View

7.2.1.2 Physical View

The *Physical* view shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself, followed by the controller and the backplane. One or more controllers are installed in the system. The controller label identifies the MegaRAID controller, such as the MegaRAID SAS 9260-8i controller, so that you can easily differentiate between multiple controllers. Each controller has one or more ports. Drives and other devices are attached to the ports. The properties for each item appear in the right panel of the screen.

Figure 176 shows the Physical view.

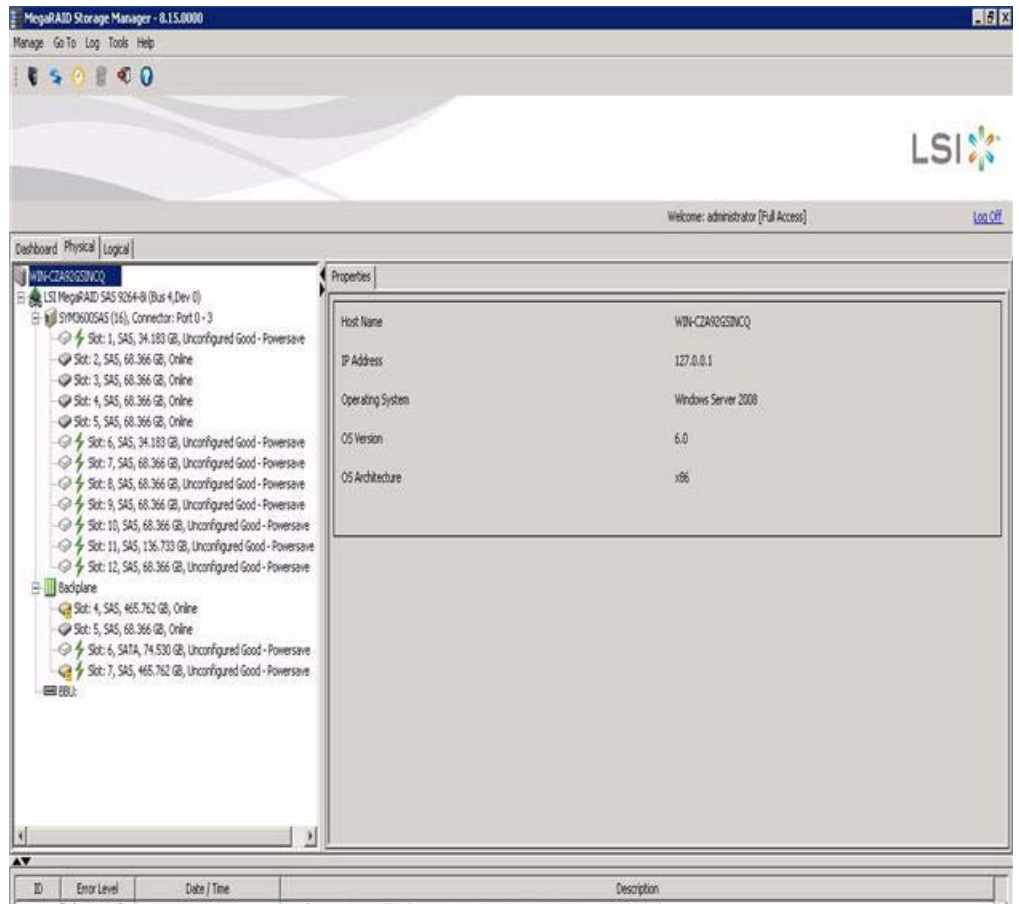


Figure 176: MegaRAID Storage Manager Physical View

7.2.1.3 Logical View

The *Logical* view shows the hierarchy of controllers, virtual drives, and the drives and drive groups that make up the virtual drives. The properties for these components appear in the right panel.

Figure 177 shows the Logical view.

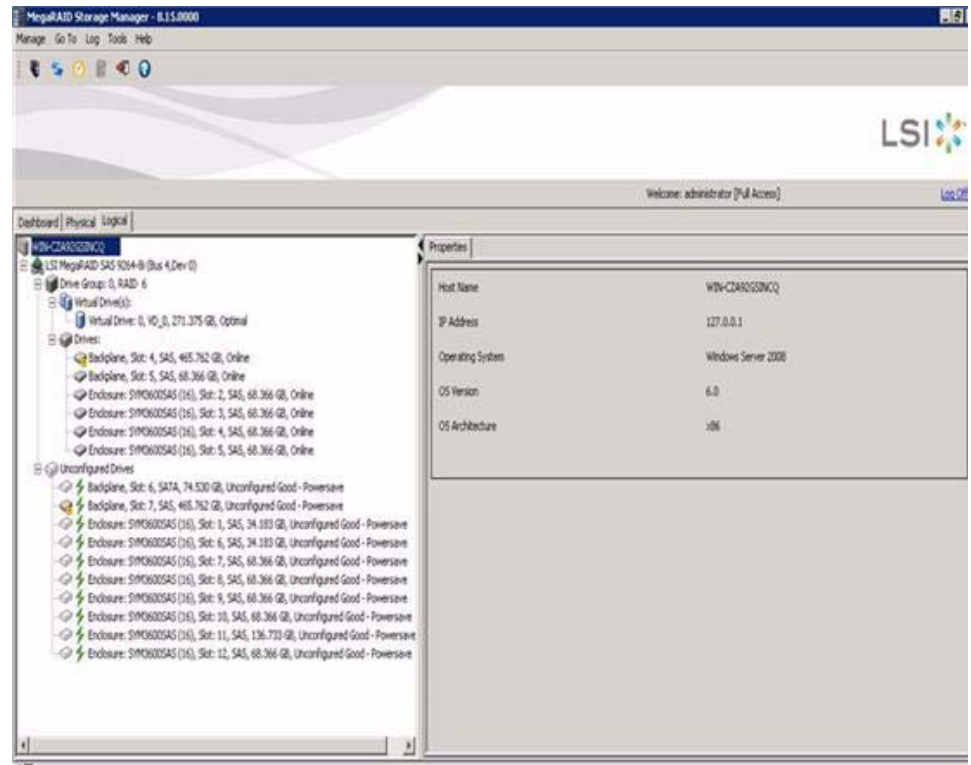


Figure 177: MegaRAID Storage Manager Logical View

7.2.1.4 Controller Properties

In the Physical view and the Logical view, you can view the chip temperature and controller temperature under the controller properties for the controller as shown in the following figure.

You can see the temperature values only for PERC controllers from PERC 7.5 onwards.

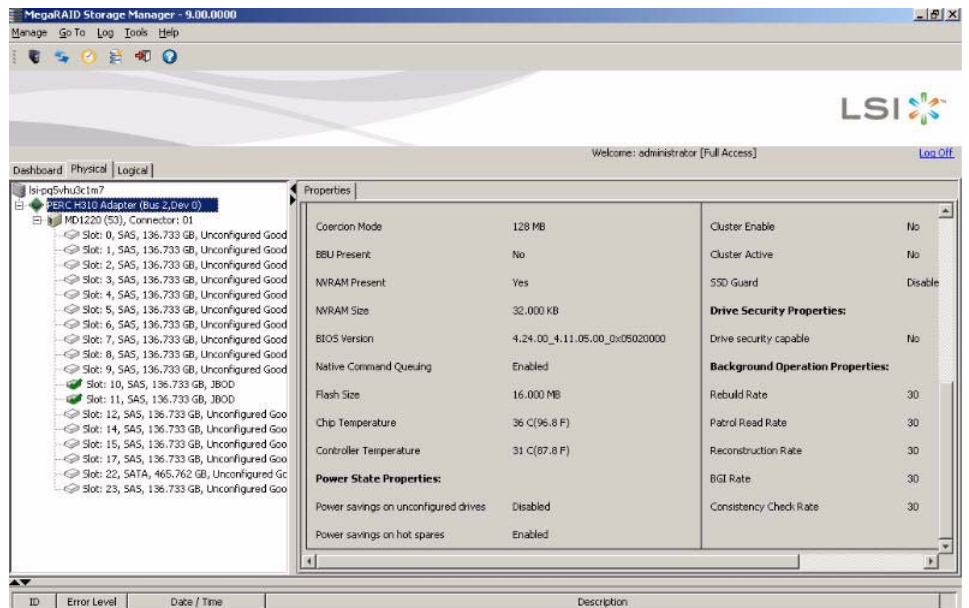


Figure 178: Chip and Controller Temperature

7.2.1.5 Physical Drive Temperature

The temperature for the physical drive is displayed in the following figure.

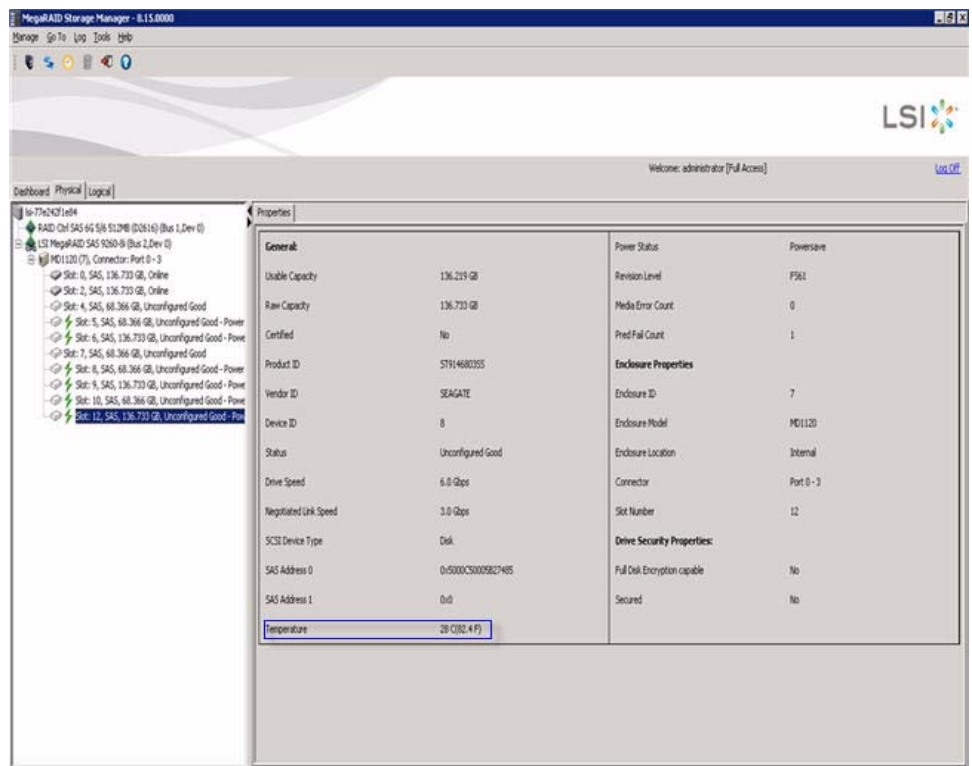


Figure 179: Physical Drive Temperature

7.2.2 Shield State

This section describes the Shield state in the MegaRAID Storage Manager software.

Physical devices in MegaRAID firmware transit between different states. If firmware detects a problem or a communication loss for a physical drive, it transitions the physical drive to a bad (FAILED/UNCONF BAD) state. To avoid transient failures, an interim state called the Shield state appears before marking the physical drive as bad state.

The Shield state is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostics tests fail, the physical drive will transition to BAD state (FAILED or UNCONF BAD).

The three possible Shield states are **Unconfigured - Shielded**, **Configured - Shielded**, and **Hotspare - Shielded**.

7.2.3 Shield State Physical View

Follow these steps to view the Shield state under the Physical view tab.

1. Click the **Physical** tab in the device tree.

The  icon indicates a Shield state.

The Physical View shield state is shown in the following figure.

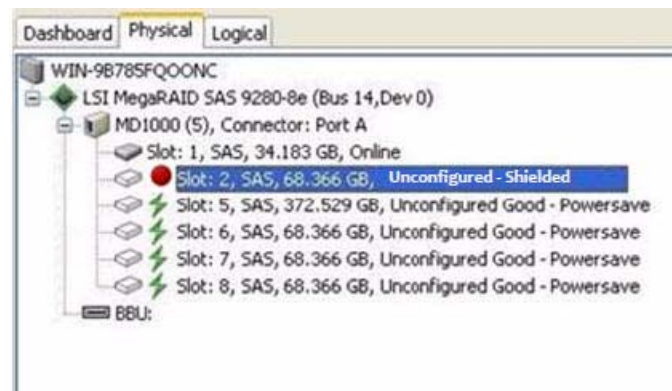



Figure 180: Physical View Shield State

7.2.4 Logical View Shield State

Follow these steps to view the Shield state under the **Logical** tab.

1. Click the **Logical** tab in the device tree.
2. The  icon indicates a Shield state.

The Logical view Shield state is shown in the following figure.

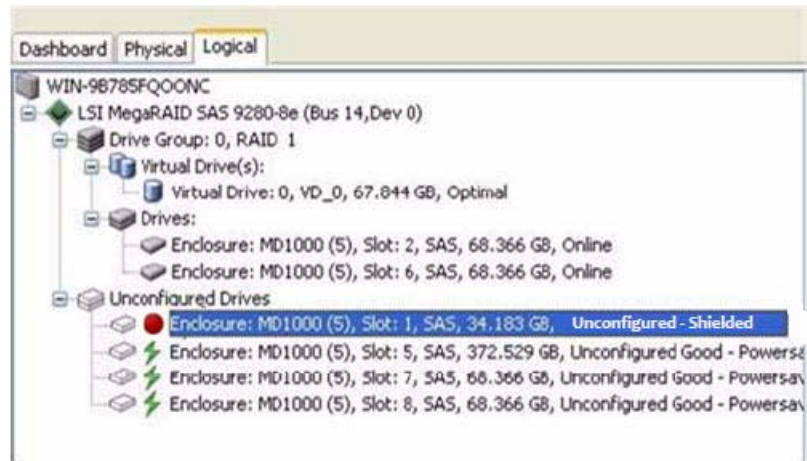



Figure 181: Logical View Shield State

7.2.5 Viewing the Physical Drive Properties

Follow these steps to view the physical properties of the drive in the Shield state.

1. Click the **Physical** tab or **Logical** tab in the device tree.
The  icon indicates a Shield state.
2. Click the physical drive in Shield state on Physical view or Logical view of the device tree to view the properties.

The device properties are displayed as shown in Figure 182.

General		Temperature	
Usable Capacity	67.044 GB	Temperature	31 C(87.8 F)
Raw Capacity	68.366 GB	Power Status	Powersave
Certified	No	Revision Level	A48B
Product ID	HU5151473VLS300	Media Error Count	0
Vendor ID	HITACHI	Pred Fail Count	0
Device ID	14	Enclosure Properties	
Status	Unconfigured - Shielded	Enclosure ID	5
Drive Speed	3.0 Gbps	Enclosure Model	MD1000
Negotiated Link Speed	3.0 Gbps	Enclosure Location	External
SCSI Device Type	Disk	Connector	Port A
SAS Address 0	0x5000CCA0074D2718	Slot Number	2
SAS Address 1	0x0	Drive Security Properties:	
		Full Disk Encryption capable	No

Figure 182: Physical Drive Properties of a Drive in Shield State

NOTE: The Status of the drive must be of the Shield type.

7.2.6 Viewing Server Profile of a Drive in Shield State

Follow these steps to view the server properties of the drive in Shield state.

1. Click the **Dashboard** tab in the device tree.
2. Click the **View Server Profile** link in the dashboard view.

The server profile information is displayed, as shown in the following figure.




Figure 183: Server Profile View of a Drive in Shield State

7.2.7 Displaying the Virtual Drive Properties

The MegaRAID Storage Manager application displays the following additional virtual drive statistics under controller properties.

- Parity size
- Mirror data size
- Metadata size

7.2.7.1 Parity Size

Parity size is used for storing parity information on RAID 5, RAID 6, RAID 50, and RAID 60 virtual drives.

Follow these steps to view the Parity Size.

1. In the Logical view, click the **Virtual Drive** node.
2. For RAID 5, RAID 6, RAID 50, and RAID 60, the **Parity Size** is displayed, as shown in [Figure 184](#).




Figure 184: Parity Size

7.2.7.2 Mirror Data Size

Mirror Data Size is used to determine the size used for storing redundant information on RAID 1 and RAID 10 virtual drives.

Follow these steps to view the **Mirror Data Size**.

1. In the **Logical** view, click on the Virtual Drive node.

The Mirror data size is displayed for RAID 1 and RAID 10 volumes, as shown in the following figure.

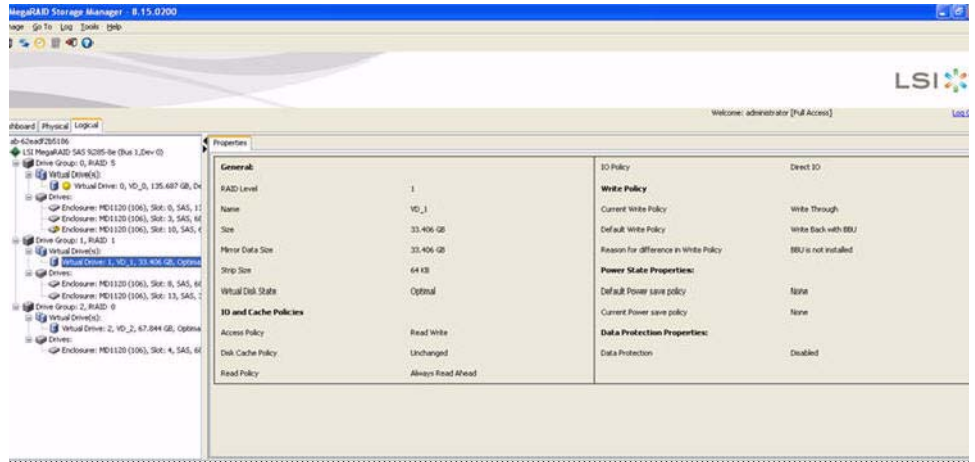


Figure 185: Mirror Data Size

NOTE: The parity size and mirror data size are not displayed for RAID 0 and RAID 00 volumes.

7.2.7.3 Metadata Size

The metadata size field displays the total space used for metadata.

Follow these steps to view the Metadata Size.

1. In the Logical view or the Physical view, click the controller node.

The total space used for metadata is displayed in this field, as shown in the following figure.



Figure 186: Metadata Size

NOTE: The size units displayed are as follows:

If the size is less than 1 MB (1024 KB), the size is displayed in KB. If the size is greater than or equal to 1 MB but less than 1 GB (1024 MB), the size is displayed in MB. If the size is greater than or equal to 1 GB, but less than 1 TB (1024 GB), the size is displayed in GB.

7.2.8 Emergency HotSpare Support

When a drive within a redundant virtual drive fails or is removed, the MegaRAID firmware automatically rebuilds the redundancy of the virtual drive by providing an Emergency Hotspare (EHSP) drive, even if no commissionable dedicated or global hot spare drive is present.

7.2.8.1 Emergency Hotspare for Physical Drives

The Emergency Hotspare property determines whether a particular drive is capable of becoming an Emergency Hotspare. This property is displayed under the controller properties only if the Global Hotspare for Emergency, and the Unconfigured Good for Emergency controller properties are enabled.

Follow these steps to view the Emergency Hotspare property.

1. In the Logical view or the Physical view, click the drive that you want to view hot spare properties.
2. The Emergency Hotspare is displayed under general properties. This property denotes whether a particular drive is capable of becoming a Hotspare, as shown in the following figure.



Figure 187: Emergency Hot Spare – Physical Drive Properties

7.2.8.2 Emergency Hotspare Property for Controllers

The Emergency Hotspare properties under the controller properties are configured based on enabling or disabling the following properties:

- Global Hotspare for Emergency
- Unconfigured Good for Emergency
- Emergency for SMARTer

Follow these steps to view the Emergency Hotspare property for controllers.

1. Click the controller node in the device tree.

The emergency hotspare properties are displayed, as shown in the following figure.

Properties			
Coercion Mode	None	SSD Guard	Disabled
BBU Present	Yes	SSD Disk Cache Setting	Disabled
NVRAM Present	Yes	Drive Security Properties:	
NVRAM Size	32,000 KB	Drive security capable	No
BIOS Version	5.15.00_4.11.05.00_0x05000000	Background Operation Properties:	
Native Command Queuing	Enabled	Rebuild Rate	34
Flash Size	16,000 MB	Patrol Read Rate	14
Memory Size	1,000 GB	Reconstruction Rate	52
Shield State Supported	Yes	BGI Rate	15
Power State Properties:		Consistency Check Rate	34
Power savings on unconfigured drives	Enabled	MegaRAID Recovery Properties:	
Power savings on hot spares	Disabled	MegaRAID Recovery	Enabled
Power Save Policy for Configured Drives	None	Emergency Hotspare Properties:	
Drive Standby Time	30mins	Global Hotspare for Emergency	Enabled
Firmware Properties:		Unconfigured Good for Emergency	Enabled
Firmware Package Version	22.0.1-0010	Emergency for SMARTer	Enabled

Figure 188: Emergency Hot Spare Properties – Controller Properties

7.2.8.3 Commissioned Hotspare

The commissioned hotspare is used to determine whether the online drive has a Commissioned Hotspare.

Follow these steps to check if the drive has Commissioned Hotspare.

1. Click the online physical drive node in the device tree.

The Commissioned Hotspare property is displayed, as shown in the following figure.

Properties	
General:	Power Status: On
Usable Capacity: 136,219 GB	Commissioned Hotspare: Yes/ No
Raw Capacity: 136,750 GB	Revision Level: 0002
Certified: No	Media Error Count: 0
Product ID: ST9146803SS	Pred Fail Count: 0
Vendor ID: SEAGATE	Enclosure Properties:
Device ID: 0	Enclosure ID: 7
Status: Online	Enclosure Model: MD1120
Drive Speed: 6.0 Gbps	Enclosure Location: Internal
Negotiated Link Speed: 3.0 Gbps	Connector: Port 4 - 7
SCSI Device Type: Disk	Slot Number: 2
SAS Address 0: 0x500050008B10460	Drive Security Properties:
SAS Address 1: 0x0	Full Disk Encryption capable: No
Temperature: 28 C (82.4 F)	Secured: No

Figure 189: Commissioned Hotspare

NOTE: This property is displayed only for online physical drives.

7.2.9 SSD Disk Cache Policy

The MegaRAID firmware provides support to change the write-cache policy for SSD media of individual physical drives.

The MegaRAID firmware does not allow any user application to modify the write-cache policies of any SSD media. The host applications can modify this property through a new logical device (LD) addition or a LD property change. When SSDs are configured in a mixed disk group with HDDs, the Physical Device Write-Cache Policy setting of all the participating drives are changed to match the SSD cache policy setting.

Follow these steps to view the SSD cache property.

1. Click the controller node in the device tree.

The Controller Properties screen appears, as shown in the following figure.

Properties			
Host Port Count	0	Backend SAS Address 6	0x0
FRU		Backend SAS Address 7	0x0
Alarm Present	Yes	Correctable Error Count	0
Alarm Enabled	Yes	Memory uncorrectable count	0
Cache Flush Interval	4 sec	Cluster Enable	No
Coercion Mode	None	Cluster Active	No
BBU Present	No	SSD Guard	Enabled
NVRAM Present	Yes	SSD Disk Cache Setting	Disabled
NVRAM Size	32,000 KB	Drive Security Properties:	
BIOS Version	3.18.00_4.09.05.00_0x0416A000	Drive security enabled	No
Native Command Queuing	Enabled	Drive security method	FDE Only
Flash Size	8,000 MB	Drive security capable	Yes
Memory Size	256,000 MB	EKM Supported	Yes
Power State Properties:		Key Management Mode	N/A
Power savings on unconfigured drives	Enabled	Background Operation Properties:	
Power savings on hot spares	Enabled	Rebuild Rate	30
Drive Standby Time	30mins	Patrol Read Rate	30
Firmware Properties:		Reconstruction Rate	30
Firmware Package Version	12.10.0-0015	BGI Rate	30
		Consistency Check Rate	30

Figure 190: Controller Properties – SSD Disk Cache Policy

7.2.9.1 Virtual Drive Settings

If the SSD cache property is enabled in the controller properties screen as shown, in [Figure 190](#), then you cannot select the disk cache policy for the virtual drives having only SSD drives or a mix of SSD drives and HDD drives during virtual drive creation. The value of the disk cache policy is unchanged and the drop-down menu is disabled.

Follow these steps to view the **Virtual Drive Settings**.

1. Right-click the controller node in the device tree.
2. Select the **Create Virtual Drive** menu option.
3. Select **Advanced Configuration**, and click **Next**.
4. Create **Drive Group**, and click **Next**.

The Create Virtual Drive – Virtual drive settings dialog appears, as shown in the following figure.

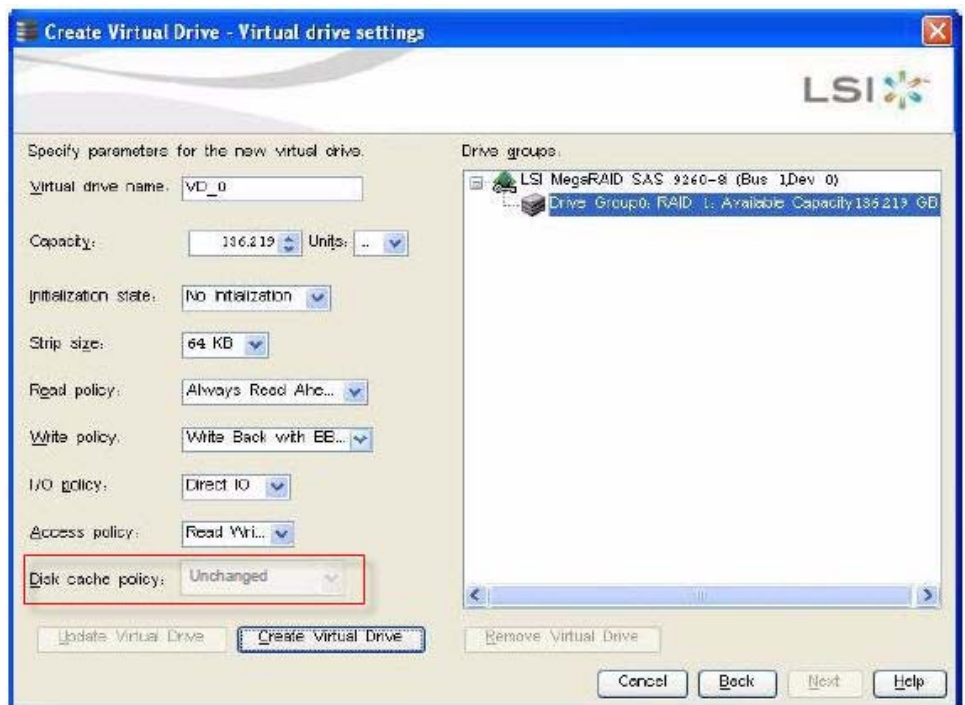


Figure 191: Virtual Drive Settings

The value of the disk cache policy is unchanged, and the drop-down list is disabled.

7.2.9.2 Set Virtual Drive Properties

Follow these steps to set virtual drive properties.

1. Right click on virtual drive node in the logical view of the device tree.
2. Select **Set Virtual Drive Properties**.

The Set Virtual Drive Properties dialog appears, as shown in [Figure 192](#).

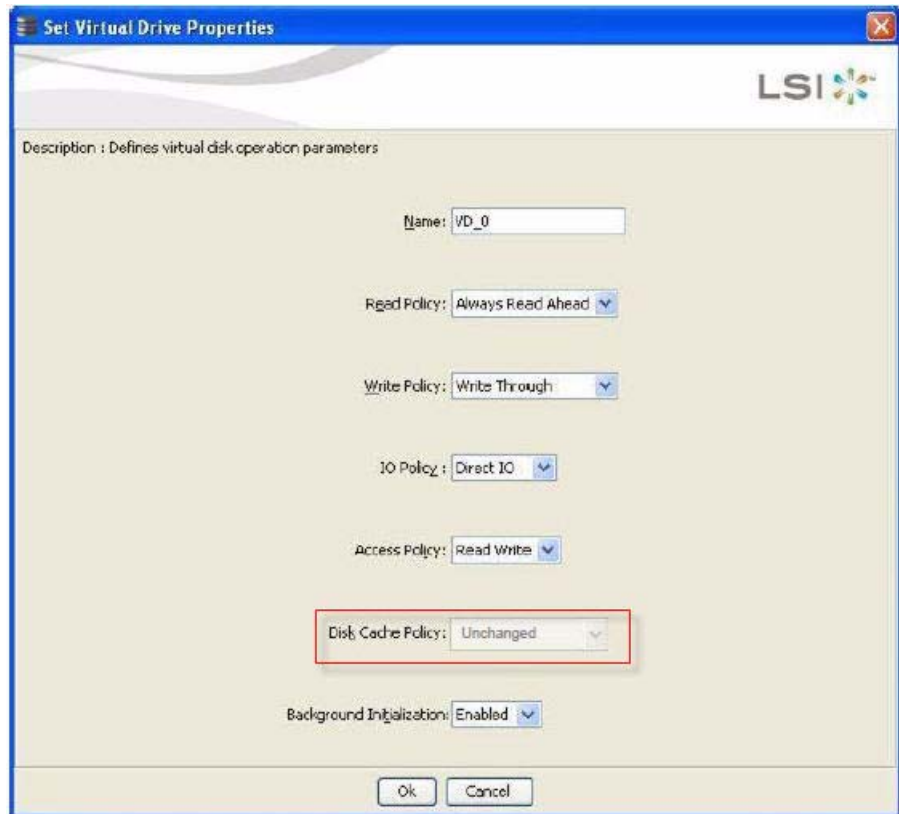


Figure 192: Virtual Drive Properties

NOTE: You cannot select the Disk cache policy for the virtual drives having only SSD drives or a mix of SSD and HDD during VD creation. The value of the Disk Cache Policy is Unchanged and can be set for only HDD drives.

7.2.10 Non-SED Secure Erase Support

This section describes the firmware changes required to securely erase data on non-SEDs (normal HDDs).

SEDs securely erase their internal encryption keys, effectively destroying all of the data present on the drive. For Non-SED drives, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The sanitization technique is more secure than a simple format operation and is commonly called a “clearing” operation, similar to the existing physical drive clear command.

Follow these steps to set physical drive properties.

1. In the Physical view, right click the **Physical Drive** node.
2. Select the **Drive Erase** option (**Alt+E**).

The **Mode Selection - Drive Erase** dialog box appears, as shown in the following figure.

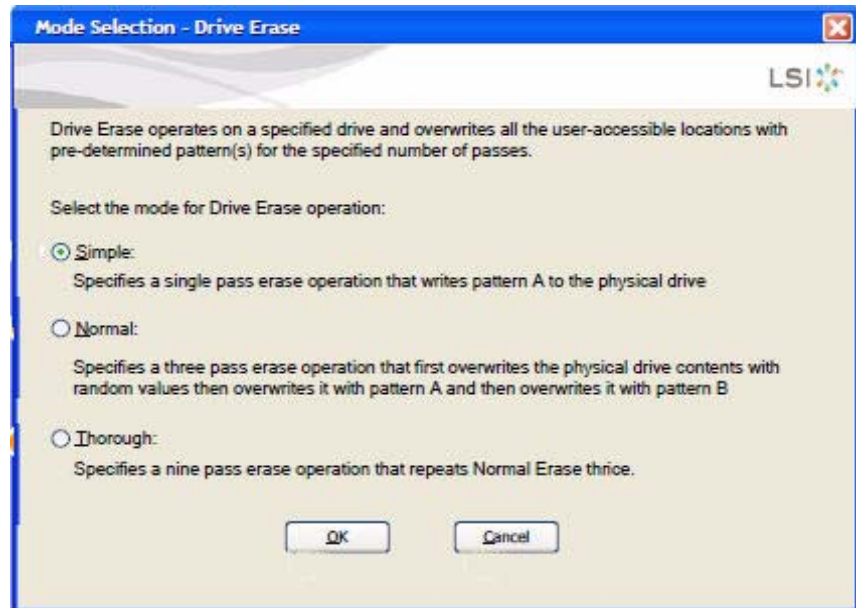


Figure 193: Mode Selection - Drive Erase Window

3. You can select the various modes available under the **Select the mode for Drive Erase operation**.

- **Simple** – (Alt + S)

When you select this option and click **OK**, the Drive Erase message box appears, as shown in the following figure.

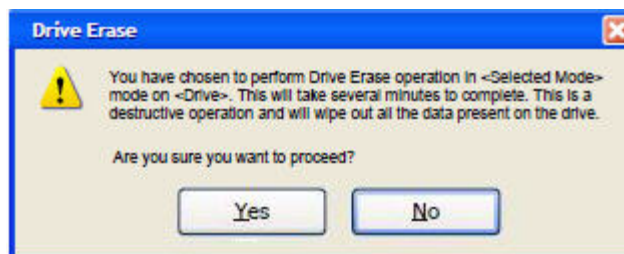


Figure 194: Drive Erase Message

- **Normal** – (Alt + N)

Select this option and click **OK**. [Figure 194](#) is displayed.

- **Thorough** – (Alt + T)

Select this option and click **OK**. [Figure 194](#) is displayed.

7.2.10.1 Group Show Progress

Physical drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

Follow these steps to check the progress of physical drive erase operation.

1. Click the **Show Progress** toolbar icon in the MegaRAID Storage Manager. You can also select **Show Progress** from the dashboard or select **Show Progress** from the Manage menu.
2. Click the **More info** link under the Background Operations portlet.

The progress bar appears, as shown in [Figure 195](#).

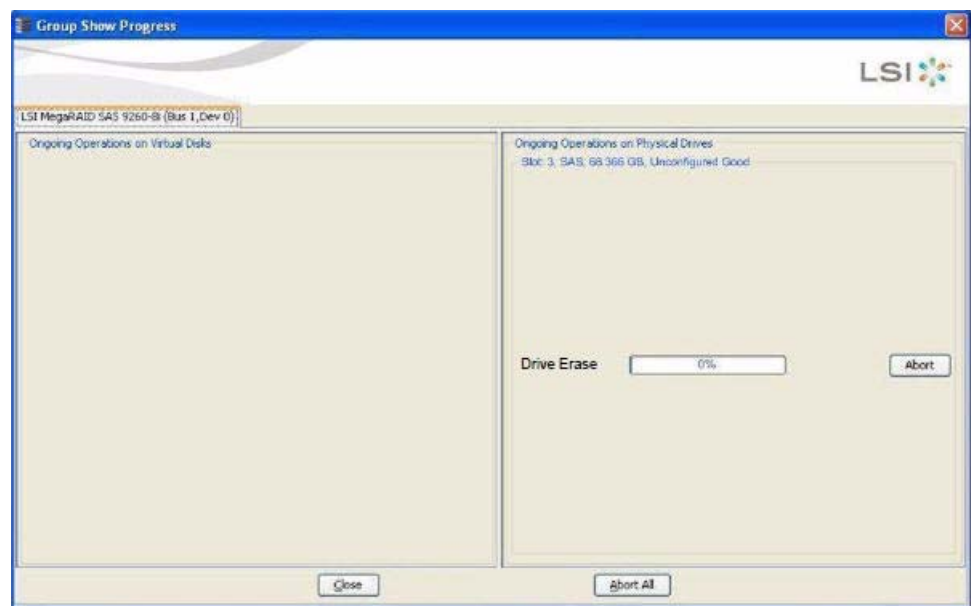


Figure 195: Group Show Progress

When you click the **Abort All** button, all Drive Erase operations stop, and the progress bar is not displayed.

7.2.10.2 Virtual Drive Erase

Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports non-zero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's LBA range. Virtual drive erase is a background operation, and it posts events to notify users of their progress.

Follow these steps to open the Virtual Drive Erase menu.

1. In the Logical view, right-click the Virtual Drive node.
2. Click on the Virtual Drive node, select top level navigation and click **Go to**.
3. Select **Virtual Drive** and select **Events & Response**.

The Logical View - Virtual Drive Erase menu appears.

4. Select **Virtual Drive Erase**.

The Virtual Drive Erase Menu opens, as shown in [Figure 196](#).

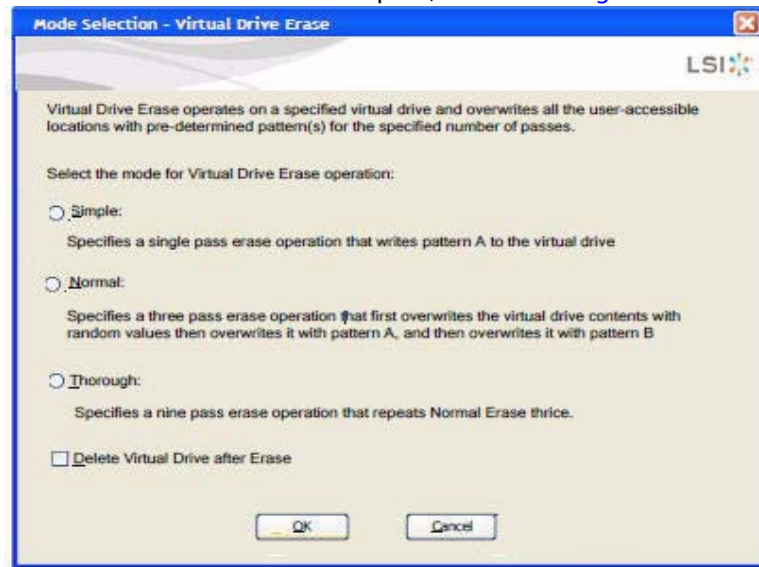


Figure 196: Mode Selection – Virtual Drive Erase Dialog

5. The menu has the following options.

- **Simple – (Alt + S)** – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 197](#) is displayed; otherwise, [Figure 198](#) is displayed.
- **Normal – (Alt + N)** – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 197](#) is displayed; otherwise, [Figure 198](#) is displayed.
- **Thorough – (Alt + T)** – After you select this option and click **OK** and if **Delete Virtual Drive after Erase** is selected, [Figure 197](#) is displayed; otherwise, [Figure 198](#) is displayed.
- **Delete Virtual Drive after Erase – (Alt + D)** – When you select this option, the virtual drive is erased and [Figure 197](#) is displayed; otherwise, [Figure 198](#) is displayed.
- **OK – (Alt + O)** – Click **OK** and if **Delete Virtual Drive after Erase** is checked, [Figure 197](#) is displayed; otherwise, [Figure 198](#) is displayed.
- **Cancel – (Alt + C)** – When you select this option, the dialog closes, and the MegaRAID Storage Manager navigates back to Physical view.

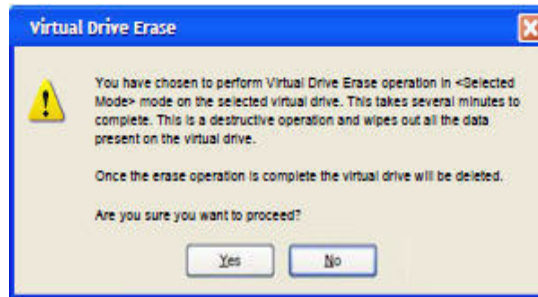


Figure 197: Warning Message for Virtual Drive Erase

- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.




Figure 198: Warning Message for Virtual Drive Erase Without Virtual Drive Delete

- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.

7.2.10.3 Group Show Progress for Virtual Drive Erase

The virtual drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

To view the progress of Group Show Progress-Virtual Drive, click the **Show Progress** toolbar icon.

You can also either select **Show Progress** from the Manage menu, or select the **More info** Link under Background Operations portlet on the dashboard.

The Virtual Drive Erase progress bar appears, as shown in [Figure 199](#).

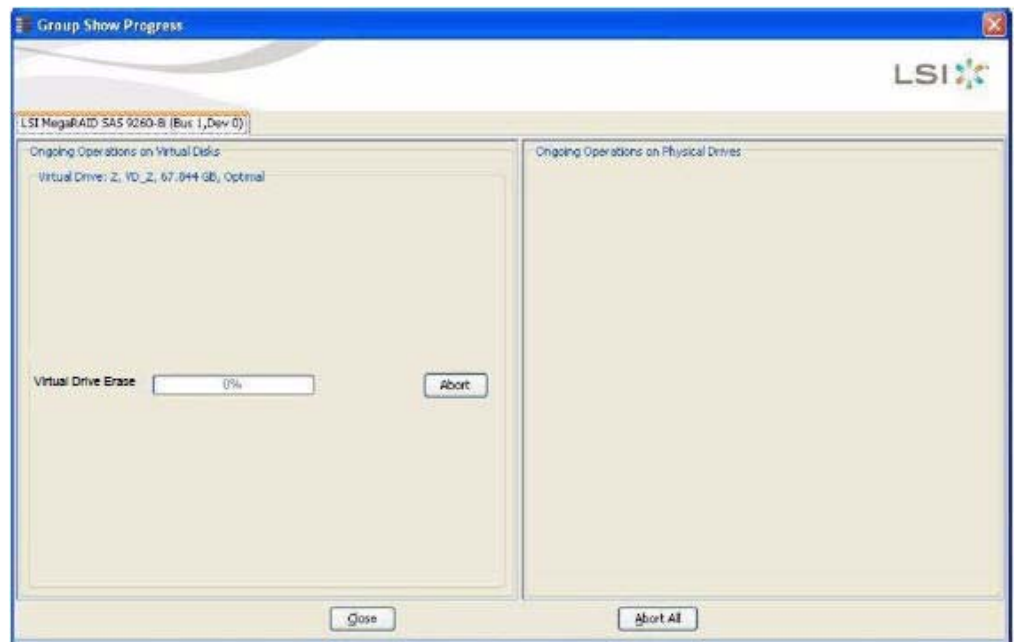


Figure 199: Group Show Progress – Virtual Drive

7.2.11 Rebuild Write Cache

MegaRAID firmware supports drive cache properties during a rebuild operation. The MegaRAID solution temporarily enables drive cache for the physical drive that is being rebuilt for the duration of the rebuild operation. Users can enable or disable this feature using the Mega CLI feature.

The MegaRAID software automatically changes the setting for a drive that is being rebuilt. If the PD_CACHE for the rebuilt drive is already set, the firmware does not need to do anything extra.

The firmware identifies and sets the cache policy of the drives whenever a rebuild operation starts and the cache policy is reflected in the event logs. The firmware also makes sure to flush the cache just before committing the drive to the disk group.

7.2.12 Background Suspend or Resume Support

MegaRAID provides a background Suspend or Resume Support feature that enhances the functionality where in the background operations running on a physical drive or a virtual drive can be suspended for some time, and resumed later using the Resume option.

The background operations, including consistency-check, rebuild, background initialization, and patrol read, are supported by an abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

A suspended operation can be resumed later by using the Resume option, and the suspended operation resumes from the point where the operation was suspended last.

To perform a suspend and resume operation, go to the Group Show Progress dialog, and perform the tasks mentioned below. You also can select **Show Progress** from the Manage menu, or select the **More info** link under the Background Operations portlet on the dashboard

The Group Show Progress dialog appears, as shown in the following figure.

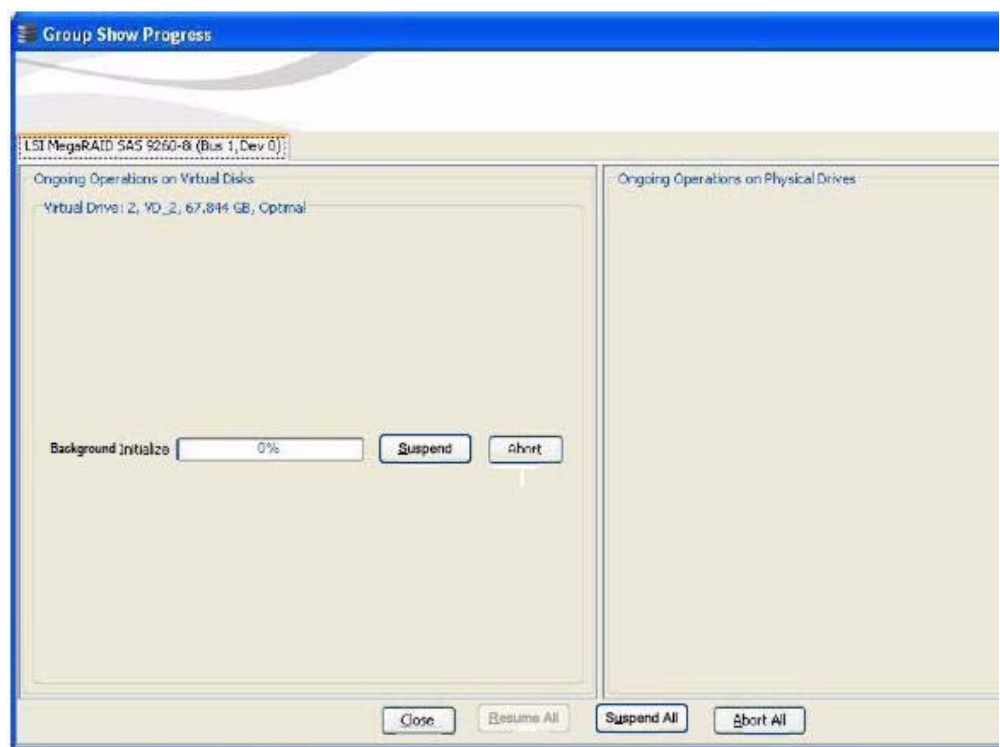



Figure 200: Group Show Progress Dialog

- **Suspend (Alt + S)** – Click the **Suspend** button to suspend the background operation taking place at that particular point of time. When the operations gets suspended, the **Resume** button appears instead of the **Suspend** button.
- **Resume (Alt + E)** – Click the **Resume** button to resume the operation from the point where it was suspended last.
- **Abort (Alt + B)** – Click the **Abort** button to abort the ongoing active operation.

- **Resume All (Alt + R)** – Click the **Resume All** button to resume all the suspended operations from the point they were suspended. This button is disabled if no operations are suspended.
- **Suspend All (Alt + S)** – Click the **Suspend All** button to suspend all the active operations. The **Suspend All** button is enabled only if one or more operations are in active state.
- **Abort All (Alt + A)** – Click the **Abort All** button to abort all the active operations.
- **Close (Alt + C)** – Click the **Close** button to close the dialog.

NOTE: The Suspend, Resume, Suspend All, and Resume All will be applicable only for background initialization and rebuild patrol read operations.

7.2.13 Enclosure Properties


To view the enclosure properties, in the Physical View click the **Enclosure**  node.

The Enclosure Properties are displayed, as shown in [Figure 201](#).

Vendor ID	DELL	FRU Number	42R5133
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	MD1000	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16VTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

Figure 201: Enclosure Properties

7.3 Monitoring Battery Backup Units

When the MegaRAID Storage Manager software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this . If it fails, a red dot appears next to the icon.

To show the properties for a BBU, perform the following steps.

1. On the main menu screen, click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.

The BBU properties appear in the right panel. The BBU properties include the following:

- The number of times the BBU has been recharged (cycle count).
- The full capacity of the BBU, plus the percentage of its current state of charge, and the estimated time until it will be depleted.
- The current BBU temperature, voltage, current, and remaining capacity.
- If the battery is charging, the estimated time until it is fully charged.
- The battery state, which says if it is in operational state.
- If battery replacement is required.
- The BBU retention time, which gives the total number of hours the battery can support the current capacity reserve.

The BBU Properties are displayed, as shown in [Figure 202](#) and [Figure 203](#).




















Properties			
BBU Battery Type	iBBU	Cycle Count	32
Battery State	Operational	Automatic Learn Cycle	Enabled
Battery Replacement	Required / Not required	Auto Learn Period	30 days
Temperature	29.0 C (84.2 F) - Normal	Next Learn Cycle	Aug 10 2010 20:52:13
Voltage	4055 mV	Relative State of Charge	99%
Current	0 mA	Absolute State of Charge	35%
Full Capacity	<value> mAh	Run Time to Empty	Battery is not being discharged
Remaining Capacity	<value> mAh	Average Time to Empty	Battery is not being discharged
BBU Retention Time	48+ Hours	Average Time to Full	Battery is not being discharged
Estimated Time to Recharge	<value> Mins	Maximum Error Margin	25%
FRU	None		

Figure 202: Battery Backup Unit Properties for iBBU Battery


Properties			
BBU Battery Type	TMM-C (Not activated) 2	Estimated Time to Recharge	<value> Mins
Battery State	Operational	FRU	None
Battery Replacement	Required / Not required	Memory Module FRU	<value>
Temperature	29.0 C (84.2 F) - Normal	Automatic Learn Cycle	Enabled
Voltage	4055 mV	Auto Learn Period	30 days
Current	0 mA	Next Learn Cycle	Aug 10 2010 20:52:13
Full Capacity	<value> Joules		
Remaining Capacity	<value> Joules		
BBU Retention Time	48+ Hours		


Figure 203: Battery Backup Unit Properties for TMM-C Battery

The following icons in the left panel represent the controllers, drives, and other devices:

	Status
	System
	Controller
	Backplane
	Enclosure
	Port
	Drive group
	Virtual drive
	Online drive
	Power save mode
	Dedicated hotspare
	Global hotspare
	Battery backup unit (BBU)
	Tape drive
	CD-ROM
	Foreign drive
	Unconfigured drive
	Locked SED
	Unlocked SED

NOTE: The MegaRAID Storage Manager software shows the icons for tape drive devices; however, no tape-related operations are supported by the utility. If these operations are required, use a separate backup application.

A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a drive has failed: .

A yellow circle to the right of an icon indicates that a device is running in a partially degraded state. For example, this icon indicates that a virtual drive is running in a degraded state because a controller has failed: .

An orange circle to the right of an icon indicates that a device is running in a degraded state.

7.3.1 Properties and Graphical View Tabs

The right panel of the MegaRAID Storage Manager window has one tab or two tabs, depending on which type of device you select in the left panel.

- The **Properties** tab displays information about the selected device. For example, if you select a controller icon in the left panel, the **Properties** tab lists information about the controller, such as the controller name, NVRAM size, and device port count. For more information, see [Section 9.15, Monitoring Controllers](#), [Section 9.16, Monitoring Drives](#), and [Section 9.18, Monitoring Virtual Drives](#).
- The **Graphical View** tab displays information about the temperature, fans, power supplies, and voltage sensors. To display a graphical view of a drive, click an enclosure icon in the left panel of the MegaRAID Storage Manager window, and click the **Graphical View** tab.

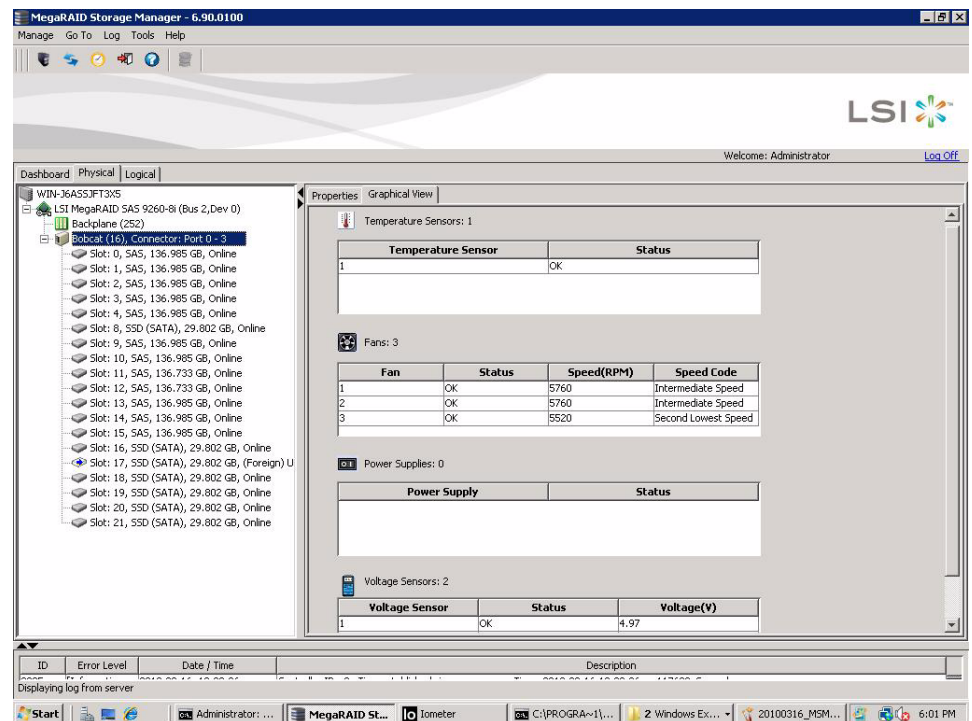


Figure 204: Properties Tab and Graphical View Tab

7.3.2 Event Log Panel

The lower part of the MegaRAID Storage Manager window displays the system event log entries. New event log entries appear during the session. Each entry has an ID, an error level indicating the severity of the event, the timestamp and date, and a brief description of the event.

For more information about the event log, see [Chapter 9, Monitoring Controllers and Its Attached Devices](#). For more information about the event log entries, see [Appendix A, Events and Messages](#).

7.3.3 Menu Bar

Here are brief descriptions of the main selections on the MegaRAID Storage Manager menu bar. Specific menu options are described in more detail in [Chapter 8](#), [Chapter 9](#), and [Chapter 10](#) of this guide.

7.3.3.1 Manage Menu

The Manage menu has a **Refresh** option for updating the display in the MegaRAID Storage Manager window (refresh is seldom required; the display usually updates automatically) and an **Exit** option to end your session on MegaRAID Storage Manager. The **Server** option shows all the servers that were discovered by a scan. In addition, you can perform a check consistency, initialize multiple virtual groups, and show the progress of group operations on virtual drives.

7.3.3.2 Go To Menu

The Go To menu is available when you select a controller, drive group, physical drive, virtual drive, or battery backup unit in the main menu screen. The menu options vary depending on the type of device selected in the left panel of the MegaRAID Storage Manager main menu. The options also vary depending on the current state of the selected device. For example, if you select an offline drive, the **Make Drive Online** option appears in the Physical Drive menu.

Configuration options are also available. This is where you access the Configuration Wizard that you use to configure drive groups and virtual drives. To access the Wizard, select the controller in the left panel, and then select **Go To >> Controller >> Create Virtual Drive**.

7.3.3.3 Log Menu

The Log menu includes options for saving and clearing the message log. For more information about the Log menu, see [Appendix A, Events and Messages](#).

7.3.3.4 Tools Menu

On the Tools menu, you can select **Tools >> Configure Alerts** to access the Configure Alerts dialog, where you can set the alert delivery rules, event severity levels, exceptions, and email settings. For more information, see [Section 9.2, Configuring Alert Notifications](#).

7.3.3.5 Help Menu

On the Help menu, you can select **Help >> Contents** to view the MegaRAID Storage Manager online help file. You can select **Help >> About MegaRAID Storage Manager** to view version information for the MegaRAID Storage Manager software.

NOTE: When you use the MegaRAID Storage Manager online help, you might see a warning message that Internet Explorer has restricted the file from showing active content. If this warning appears, click on the active content warning bar, and enable the active content.

NOTE: If you are using the Linux operating system, you must install Firefox[®] browser or Mozilla[®] browser for the MegaRAID Storage Manager online help to display.

NOTE: When connected to the VMware server, only the IP address and the host name information appear. The other information, such as the operating system name, version, and architecture do not appear.

Chapter 8

Configuration

This chapter explains how to use MegaRAID Storage Manager software to create and modify storage configurations on LSI SAS controllers.

The LSI SAS controllers support RAID 0, RAID 1, RAID 5, RAID 6, RAID 00, RAID 10, RAID 50, and RAID 60 storage configurations. The Configuration wizard allows you to easily create new storage configurations and modify the configurations. To learn more about RAID and RAID levels, see [Chapter 2, Introduction to RAID](#).

NOTE: You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

8.1 Creating a New Storage Configuration

You can use the MegaRAID Storage Manager software to create new storage configurations on systems with LSI SAS controllers. You can create the following types of configurations:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you. This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize virtual drive creation. This option provides greater flexibility when creating virtual drives for your specific requirements.

This section describes the virtual drive parameters and explains how to create simple and advanced storage configurations.

8.1.1 Selecting Virtual Drive Settings

This section describes the virtual drive settings that you can select when you use the advanced configuration procedure to create virtual drives. You should change these parameters only if you have a specific reason for doing so. It is usually best to leave them at their default settings.

- **Initialization state:** Initialization prepares the storage medium for use. Specify the initialization status:
 - **No Initialization:** (the default) The new configuration is not initialized, and the existing data on the drives is not overwritten.
 - **Fast Initialization:** The firmware quickly writes 0s to the first and last 8-MB regions of the new virtual drive and then completes the initialization in the background. This allows you to start writing data to the virtual drive immediately.
 - **Full Initialization:** A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.

NOTE: BGI is supported only for RAID 5 and RAID 6 and not for any other RAID levels. New RAID 5 virtual drives require at least five drives for a background initialization to start. New RAID 6 virtual drives require at least seven drives for a background initialization to start. If there are fewer drives, the background initialization does not start.

- **Strip size:** Strip sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB are supported. The default is 64 KB. For more information, see the *striping* entry in the Glossary.
- **Read policy:** Specify the read policy for this virtual drive:
 - **Always read ahead:** Read ahead capability 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 process speeds up reads for sequential data, but there is little improvement occurs when accessing random data.
 - **No read ahead:** (the default) Disables the read ahead capability.
- **Write policy:** Specify the write policy for this virtual drive:
 - **Write Through:** In this 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 option eliminates the risk of losing cached data in case of a power failure.
 - **Always Write Back:** In this 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.
 - **Write Back with BBU:** (the default) In this mode, the controller enables write back caching when the battery backup unit (BBU) is installed and charged. This option provides a good balance between data protection and performance.

NOTE: The write policy depends on the status of the BBU. If the BBU is not present, is low, is failed, or is being charged, the current write policy switches to write through, which provides better data protection.

- **I/O policy:** The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - **Cached IO:** In this mode, all reads are buffered in cache memory.
 - **Direct IO:** (the default) In this 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.
Cached IO provides faster processing, and **Direct IO** ensures that the cache and the host contain the same data.
- **Access policy:** Select the type of data access that is allowed for this virtual drive.
 - **Read/Write:** (the default) Allow read/write access. This setting is the default value.
 - **Read Only:** Allow read-only access.
 - **Blocked:** Do not allow access.

- **Disk cache policy:** Select a cache setting for this drive:
 - **Enabled:** Enable the disk cache.
 - **Disabled:** Disable the disk cache.
 - **Unchanged:** (the default) Leave the current disk cache policy unchanged.

8.1.2 Optimum Controller Settings for CacheCade

Write Policy: Write Back

IO Policy: Cached IO

Read Policy: No Read Ahead

Stripe Size: 64 KB

8.1.3 Optimum Controller Settings for FastPath

Write Policy: Write Through

IO Policy: Direct IO

Read Policy: No Read Ahead

Stripe Size: 64 KB

8.1.4 Creating a Virtual Drive Using Simple Configuration

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

NOTE: You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in [Section 8.1.5, Creating a Virtual Drive using Advanced Configuration](#).

Follow these steps to create a new storage configuration in simple configuration mode.

1. Perform either of the following steps:
 - Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window, and select **Create Virtual Drive**.
 - Select the controller node, and select **Go To>>Controller>>Create Virtual Drive** in the menu bar, as shown in [Figure 205](#).

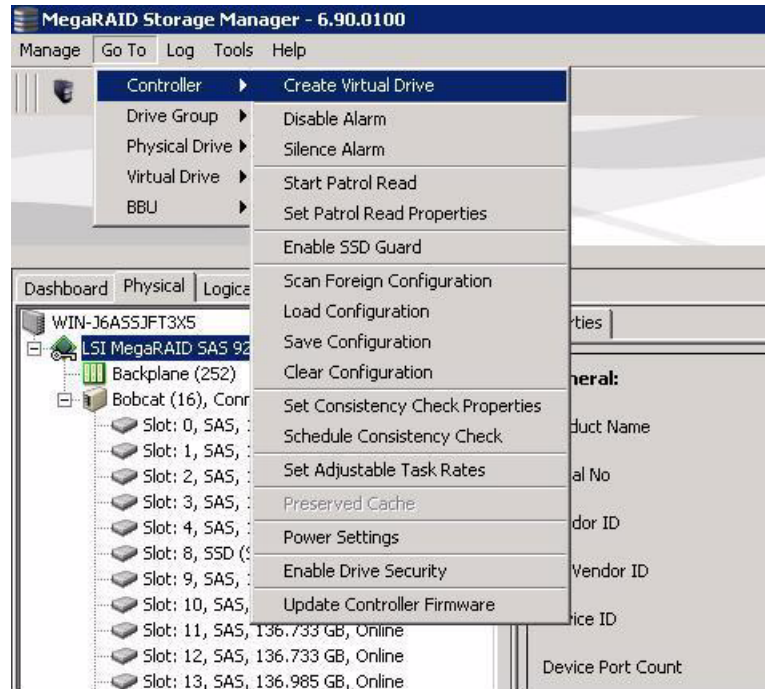


Figure 205: Create Virtual Drive Menu Option

The dialog for the configuration mode (simple or advanced) appears, as shown in the following figure.

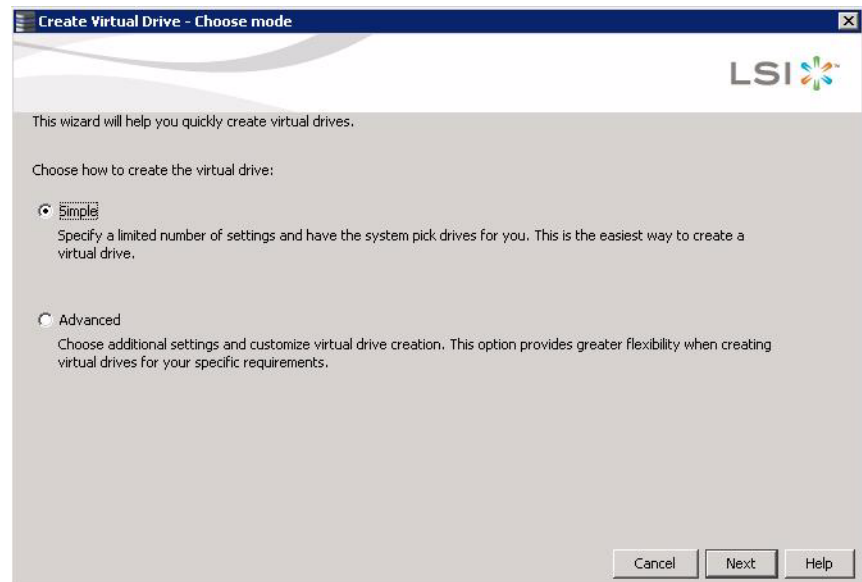


Figure 206: Create Virtual Drive – Choose Mode

2. Select the **Simple** radio button and click **Next**.

The Create Virtual Drive - Allocate capacity dialog appears, as shown in the following figure. If unconfigured drives are available, you have the option to use those unconfigured drives. If unconfigured drives are available, the Create Drive Group Settings window appears and you can go to step 4.

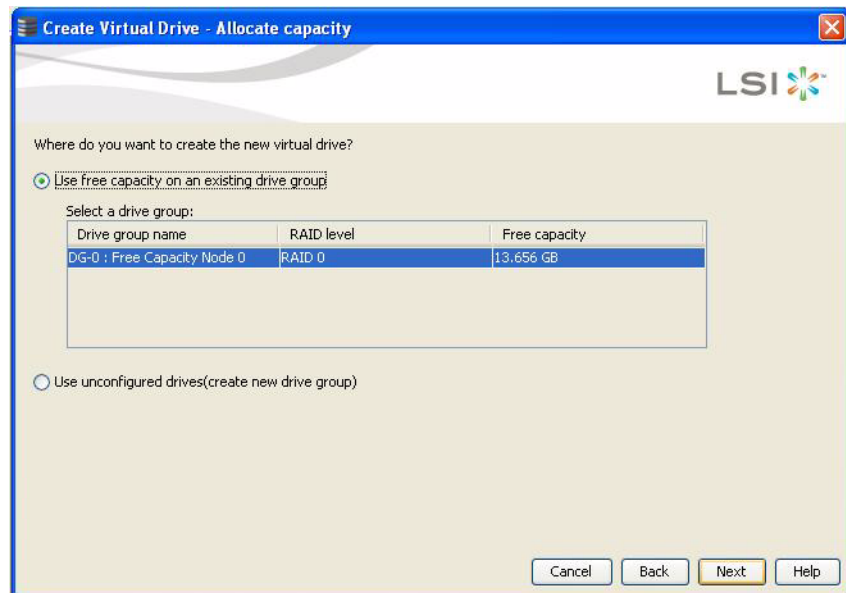


Figure 207: Using the Free Capacity of an Existing Drive Group

3. Perform either of the two options:
 - If a drive group exists, select the **Use free capacity on an existing drive group** radio button and click **Next**. Continue with step 4.

The Create Virtual Drive window appears, as shown in [Figure 208](#). If different types of drives are attached to the controller, such as HDD, SDD, SAS, and SATA, an option appears to allow drive type mixing.

- If unconfigured drives are available, select the radio button to use the unconfigured drives, and click **Next**. Continue with step 10.

The Summary window appears as shown in [figure 197](#).

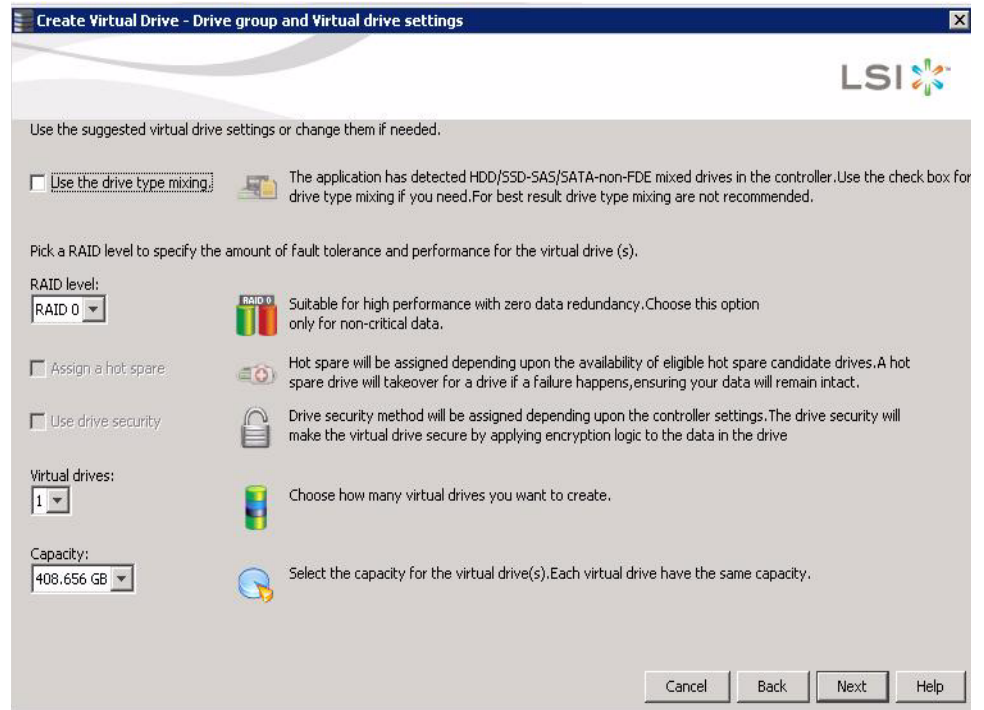


Figure 208: Create Virtual Drive – Drive group and Virtual drive settings Dialog

4. If you want to allow different types of drives in a configuration, select the **Use the drive type mixing** check box.

NOTE: For best results, do not use drive type mixing.

5. Select the RAID level desired for the virtual drive.

When you use simple configuration, the RAID controller supports RAID levels 1, 5, and 6. In addition, it supports independent drives (configured as RAID 0). The window text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available. To learn more about RAID levels, see [Chapter 2, Introduction to RAID](#).

6. Select the **Assign a hot spare** check box if you want to assign a dedicated hot spare to the new virtual drive.

If an unconfigured good drive is available, that drive is assigned as a hot spare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).

7. Select the **Use drive security** check box if you want to set a drive security method.

The LSI SafeStore Data Security Service encrypts data and provides disk-based key management for your data security solution. This solution protects the data in the event of theft or loss of drives. Refer to [Section 11.6, LSI SafeStore Encryption Services](#) for more information about the SafeStore feature.

8. Use the drop-down list in the **Virtual drives** field to choose how many virtual drives you want to create.
9. Select the capacity of the virtual drives.
Each virtual drive has the same capacity.
10. Click **Next**.

The Create Virtual Drive - Summary window appears, as shown in the following figure. This window shows the selections you made for simple configuration.

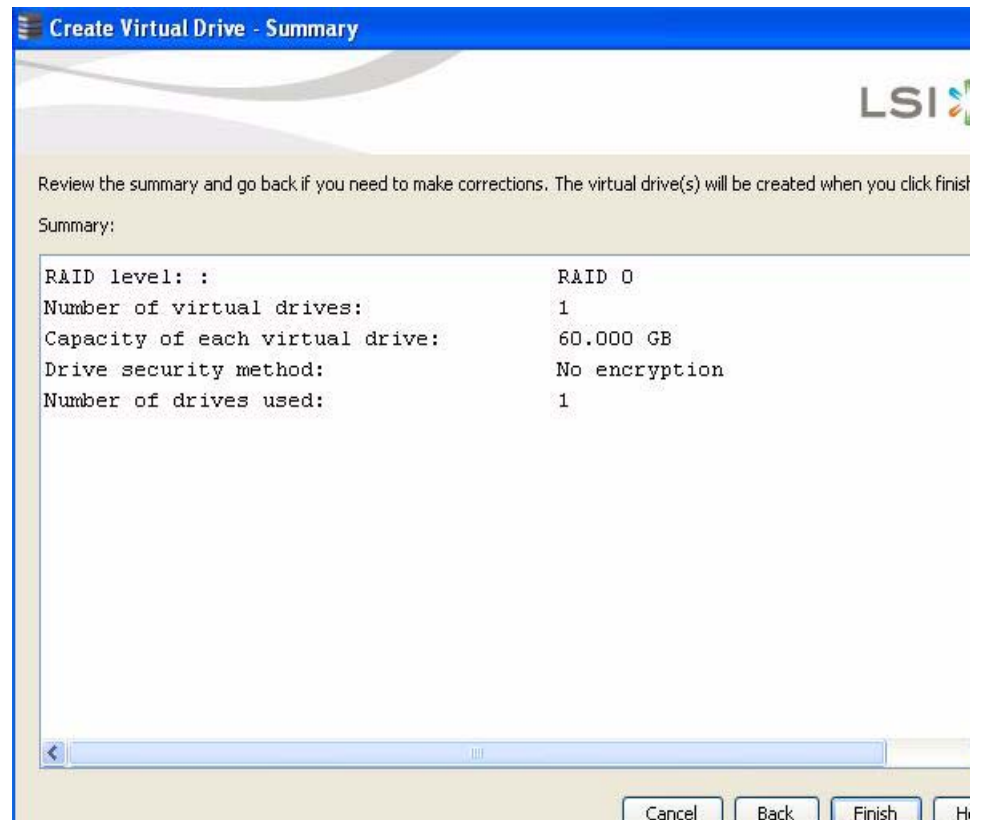


Figure 209: Create Virtual Drive - Summary Window

11. Either click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

NOTE: If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog box that identifies these drives appears.

8.1.5 Creating a Virtual Drive using Advanced Configuration

The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

Follow these steps to create a new storage configuration in the advanced configuration mode. This example shows the configuration of a spanned drive group.

1. Perform either of the following steps to bring up the Configuration wizard:
 - Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window, and select **Create Virtual Drive**.
 - Select the controller node, and select **Go To >> Controller >> Create Virtual Drive** in the menu bar, as shown in the following figure.



Figure 210: Create Virtual Drive Menu

The dialog for the choosing the configuration mode (simple or advanced) appears, as shown in the following figure.

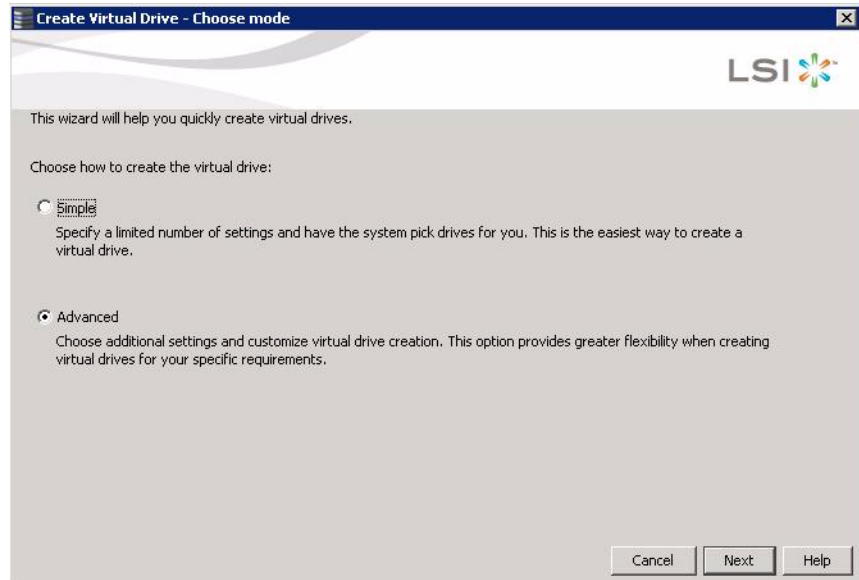


Figure 211: Create Virtual Drive - Choose mode Dialog

2. Select the **Advanced** radio button and click **Next**.

The Create Drive Group Settings window appears, as shown in the following figure.

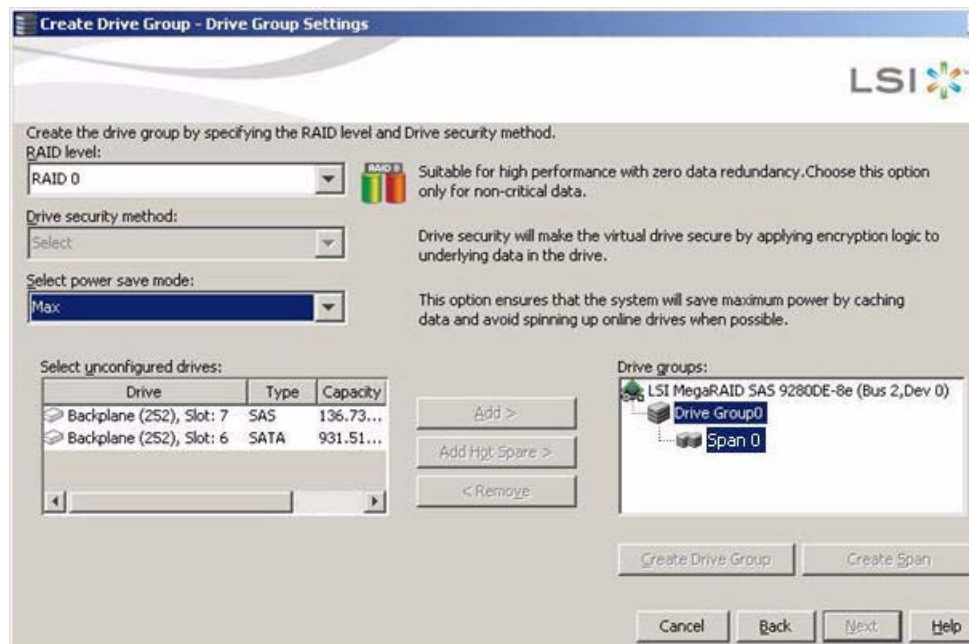


Figure 212: Create Drive Group – Drive Group Settings Window

3. Select the following items on the Create Drive Group - Drive Group Settings Window:
 - a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The dialog text gives a brief description of the RAID level that you select. You can choose the RAID levels depending on the number of available drives. To learn more about RAID levels, see [Chapter 2, Introduction to RAID](#).

- b. Scroll down the menu for the **Drive security method** field if you want to set a drive security method.

The drive security feature provides the ability to encrypt data and use disk-based key management for your data security solution. This solution provides protection to the data in the event of theft or loss of drives. See [Section 11.6, LSI SafeStore Encryption Services](#), for more information about drive security and encryption.

- c. Select *unconfigured* drives from the list of drives and click **Add>** to add them to the drive group.

The selected drives appear under **Span 0** below **Drive Group 0**, as shown in the following figure.

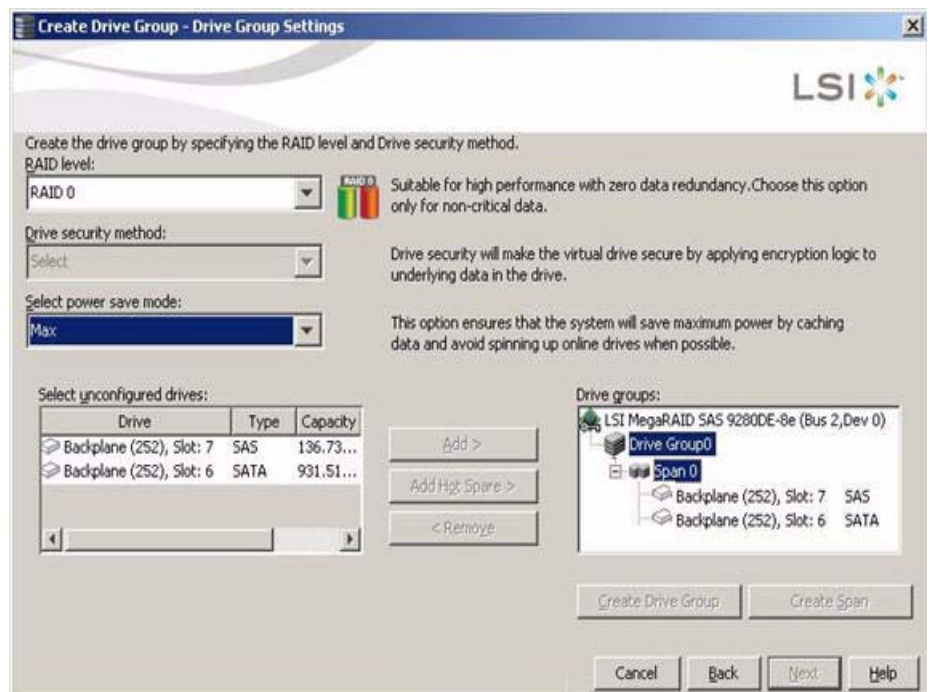


Figure 213: Span 0 of Drive Group 0

- d. Click **Create Span** to create a second span in the drive group.
- e. Select *unconfigured* drives from the list of drives and click **Add>** to add them to the second drive group.

The selected drives appear under **Span 1** below **Drive Group 0**, as shown in [Figure 214](#).

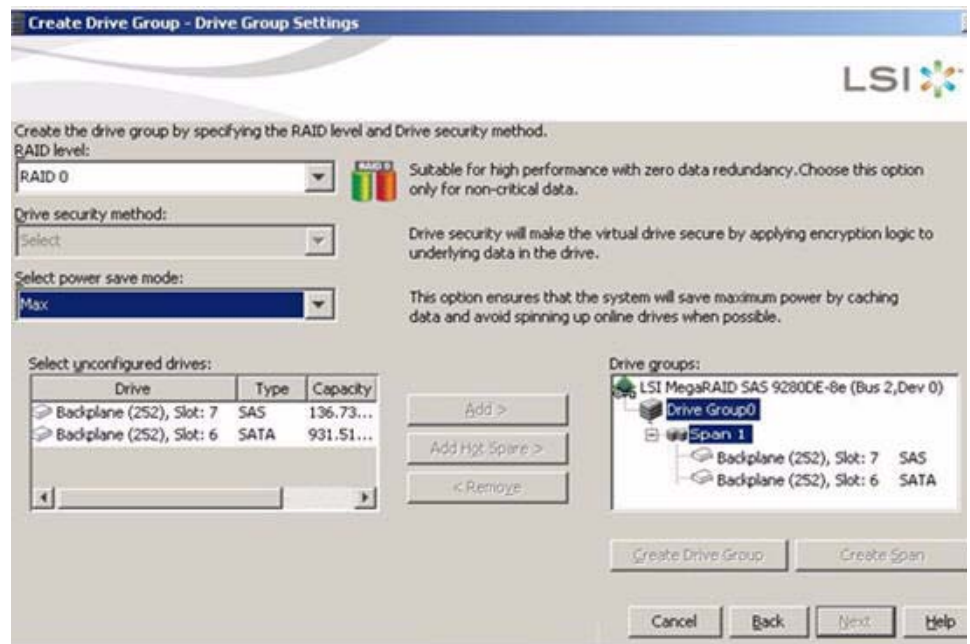


Figure 214: Span 0 and Span 1 of Drive Group 0

- f. Click **Create Drive Group** to make a drive group with the spans.
- g. Click **Next** to complete this step.

The Create Virtual Drive - Virtual drive settings window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

NOTE: The parameters in the Create Virtual Drive – Virtual drive settings window display in Disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.

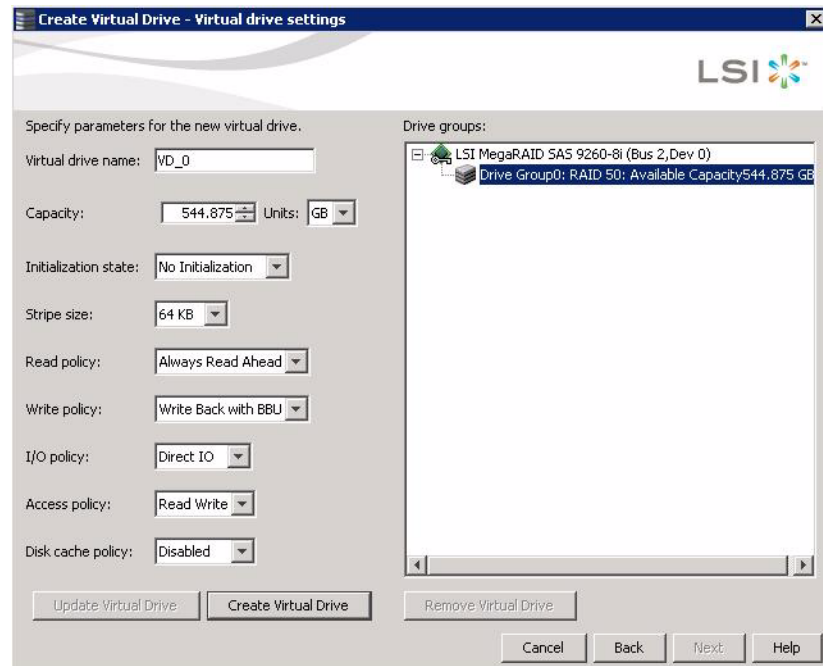


Figure 215: Create Virtual Drive -Virtual Drive Settings Window

NOTE: If you select **Write Back with BBU** as the write policy, and no battery exists, the battery is low or failed, or the battery is running through a re-learn cycle, the write policy switches to **Write Through**. This setting eliminates the risk of data loss in case of a power failure. A message window notifies you of this change.

4. Change any virtual drive settings, if desired.

See [Section 8.1.1, Selecting Virtual Drive Settings](#) for more information about the virtual drive settings.

5. Click **Create Virtual Drive**.

The new virtual drive appears under the drive group. The options **Update Virtual Drive** and **Remove Virtual Drive** are available. **Update Virtual Drive** allows you to change the virtual drive settings, and **Remove Virtual Drive** allows you to delete the virtual drive.

6. Click **Next**.

The Create Virtual Drive - Summary window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

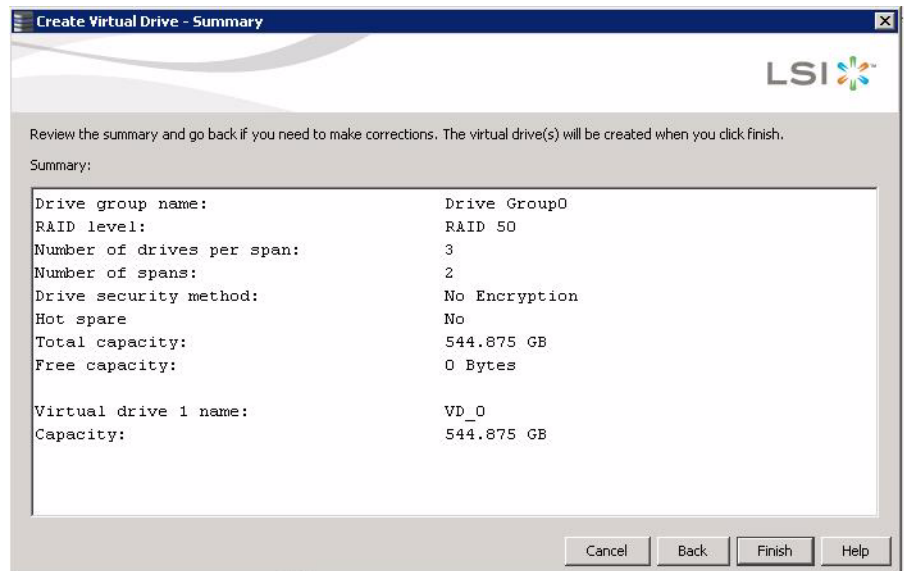


Figure 216: Create Virtual Drive – Summary Window

7. Click **Back** to return to the previous window to change any selections or click **Finish** to accept and complete the configuration.

After you click **Finish**, the new storage configuration is created and initialized according to the selected options.

NOTE: If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog appears that identifies the drives.

After the configuration is completed, a dialog notifies you that the virtual drives were created successfully.

8. Click **OK**. The Enable SSD Caching on New Virtual Drives dialog appears.

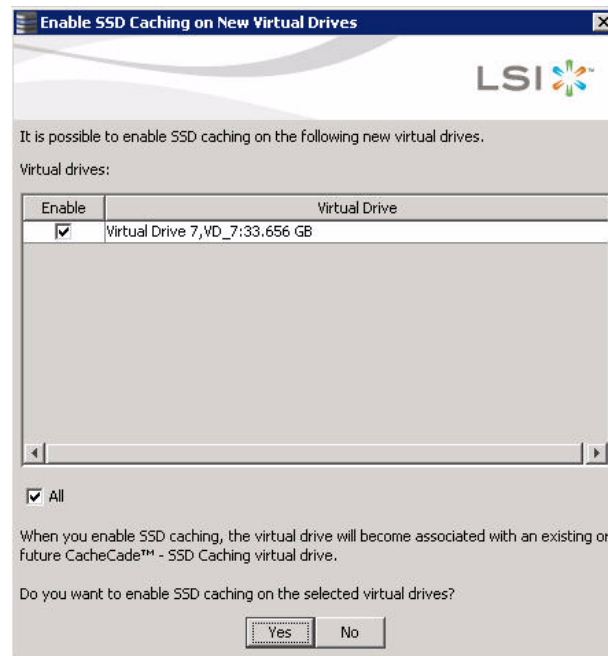


Figure 217: Enable SSD Caching on New Virtual Drives

The newly created virtual drive is enabled for SSD caching by default.

9. Click **OK** to confirm SSD caching on the virtual drive. Click **No** if you want to disable SSD caching on the virtual drive.

The **All** check box is selected by default. To disable SSD caching on the virtual drives, deselect the **All** check box.

If more drive capacity exists, the dialog asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

10. Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.

If you select **Yes**, the system takes you to the Create Virtual Drive window, as shown in [Figure 208](#). If you select **No**, the utility asks whether you want to close the wizard.

11. If you selected **No** in the previous step, select either **Yes** or **No** to indicate whether you want to close the wizard.

If you select **Yes**, the configuration wizard closes. If you select **No**, the dialog closes, and you remain on the same page.

8.2 Converting JBOD Drives to Unconfigured Good

You can convert JBOD drives to Unconfigured Good using the **Create Virtual Drive** option or **Make Unconfigured Good** drive option with a single configuration.

Perform the following steps to configure JBOD to Unconfigured Good drives by using any one of the following options.

- Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window, and select **Create Virtual Drive**.
- Select the controller node, and select **Go To>>Controller>>Create Virtual Drive**.

The **Create Virtual Drive - JBOD to Unconfigured Good Conversion** wizard appears, as shown in the following figure.

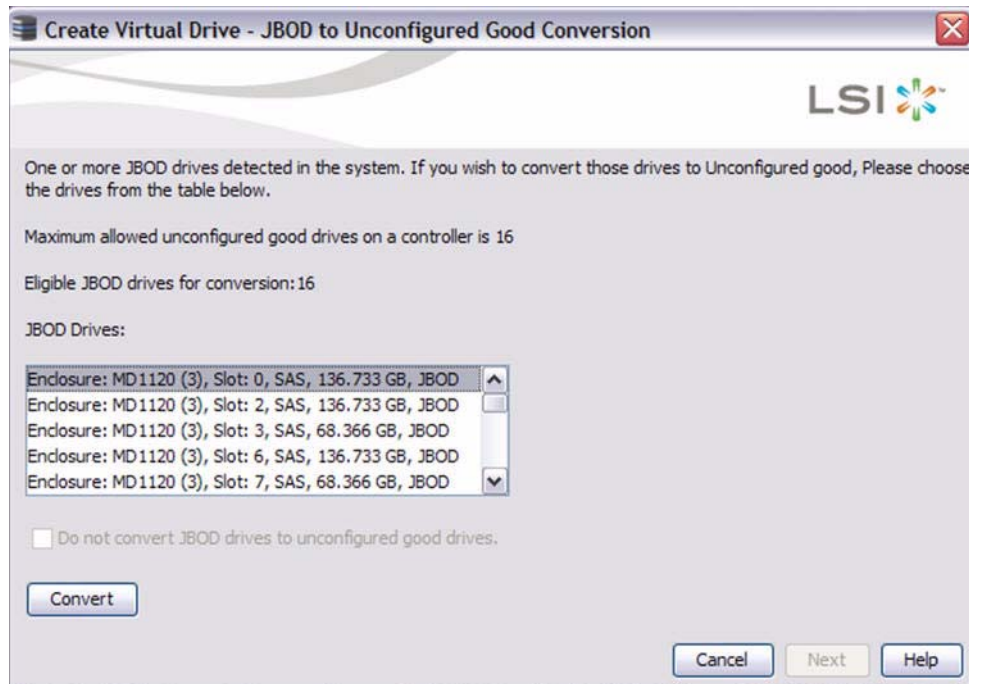


Figure 218: Create Virtual Drive - JBOD to Unconfigured Good Conversion Dialog

The **JBOD Drives** field displays the available JBOD drives available in the system.

12. Select the drives which you want configured as Unconfigured Good and then click **Convert**. Clicking on **Convert** configures the selected JBODs to Unconfigured Good Drives.

NOTE: If you do not wish to make any JBOD as unconfigured good drives, select the **Do not convert JBOD drives to unconfigured good drives** check box, and the MegaRAID Storage Manager application skips changing any selected JBOD to unconfigured good drive.

13. Click **Next**.

The **Create Virtual drive** Screen appears as shown in [Figure 208](#).

8.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Window

You can also convert JBOD to Unconfigured Good by selecting the **Controller >>Make UnConfigured Good** option from the main MegaRAID Storage Manager window.

The Make Configured Good dialog appears, as shown in the following figure.



Figure 219: Make Configured Good Dialog

14. Select the JBOD drives to be configured as unconfigured good.
15. Click **OK**.

The selected JBOD drives are configured as unconfigured good.

8.3 Adding Hot Spare Drives

Hot spares are drives that are available to automatically replace failed drives in a RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, or RAID 60 virtual drive. *Dedicated hot spares* can be used to replace failed drives in a selected drive group only. *Global hot spares* are available to any virtual drive on a specific controller.

To add a dedicated or global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.
For each drive, the window displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.
2. Either select **Go To>>Physical Drive>>Assign Global Hot Spare**, or select **Go To>>Physical Drive>>Assign Dedicated Hot Spare**.
3. If you selected **Assign Dedicated Hotspare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.
If you selected **Assign Global Hotspare**, skip this step, and go to the next step. The hot spare is available to any virtual drive on a specific controller.
4. Click **Go** to create the hot spare.

The drive state for the drive changes to dedicated or global hot spare, depending on your selection.

8.4 Changing Adjustable Task Rates

If you want to change the Rebuild rate and other task rates for a controller, you must first log onto the server in Full Access mode.

NOTE: It is LSI recommended that you leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rates below the defaults, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default values here so you can restore them later, if necessary:

Rebuild rate: _____

Background Initialization (BGI) rate: _____

Check consistency rate: _____

To change the adjustable task rates, perform the following steps:

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To>>Controller>>Set Adjustable Task Rates** from the menu bar, as shown in the following figure.

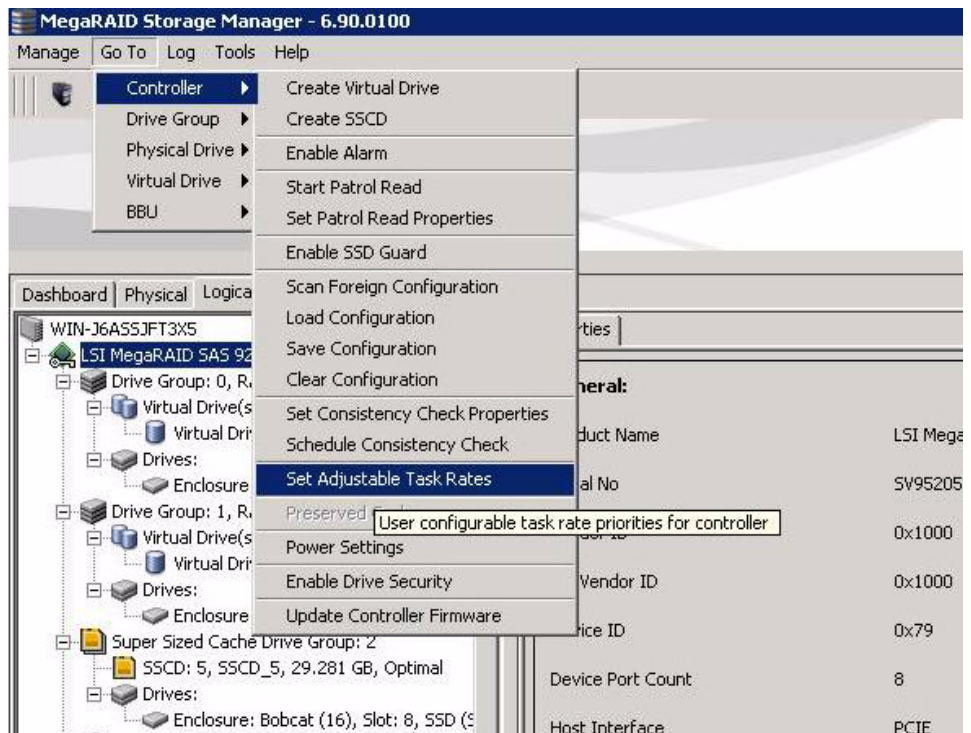


Figure 220: Set Adjustable Task Rates Menu

The Set Adjustable Task Rates window appears, as shown in the following figure.

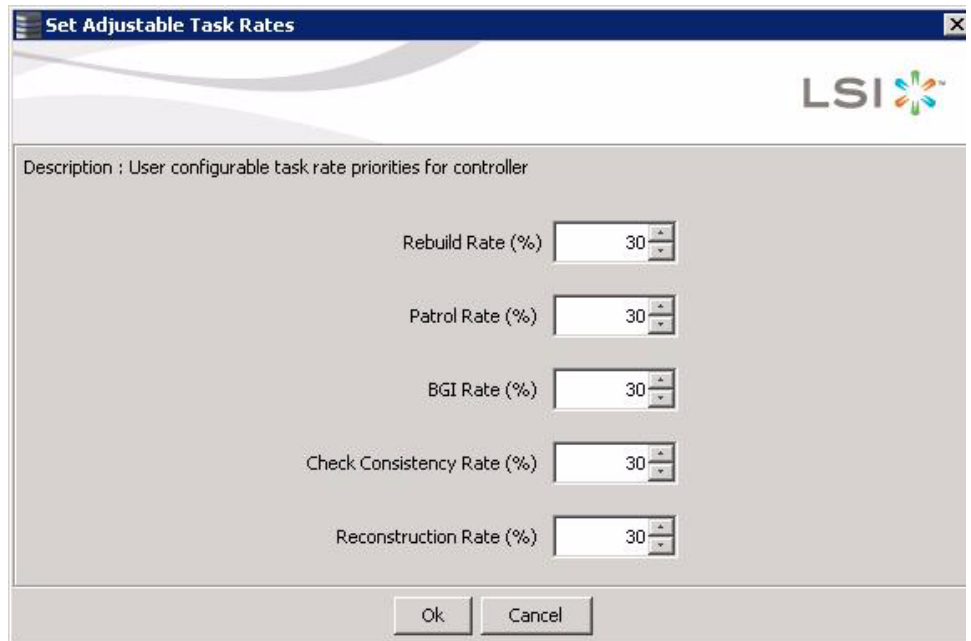


Figure 221: Set Adjustable Task Rates Menu

3. Enter changes, as needed, to the following task rates:
 - **Rebuild Rate.** Enter a number from 0 to 100 to control the rate at which a rebuild will be performed on a drive when one is necessary. The higher the number, the faster the rebuild will occur (and the system I/O rate may be slower as a result).
 - **Patrol Rate.** Enter a number from 0 to 100 to control the rate at which patrol reads will be performed. Patrol read monitors drives to find and resolve potential problems that might cause drive failure. The higher the number, the faster the patrol read will occur (and the system I/O rate may be slower as a result).
 - **Background Initialization (BGI) Rate.** Enter a number from 0 to 100 to control the rate at which virtual drives are initialized “in the background.” Background initialization establishes mirroring or parity for a RAID virtual drive while allowing full host access to the virtual drive. The higher the number, the faster the initialization will occur (and the system I/O rate may be slower as a result).
 - **Check Consistency Rate.** Enter a number from 0 to 100 to control the rate at which a consistency check is done. A consistency check scans the consistency data on a fault tolerant virtual drive to determine if the data has become corrupted. The higher the number, the faster the consistency check is performed (and the system I/O rate may be slower as a result).
 - **Reconstruction Rate.** Enter a number from 0 to 100 to control the rate at which reconstruction of a virtual drive occurs. The higher the number, the faster the reconstruction occurs (and the system I/O rate may be slower as a result).
4. Click **Ok** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

8.5 Changing Power Settings

The RAID controller includes Dimmer Switch technology that conserves energy by placing certain unused drives into Power-Save mode. In Power-Save mode, the drives use less energy, and the fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application timeouts caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

You can use the **Power Settings** field in the MegaRAID Storage Manager software to choose whether to allow unconfigured drives or Commissioned Hotspares to enter Power-Save mode.

NOTE: The Dimmer Switch technology is enabled by default.

When they are in the Power-Save mode, unconfigured drives and drives configured as Commissioned Hotspares (dedicated or global) can be spun down. When spun down, the drives stay in Power-Save mode except for periodic maintenance, which includes the following:

- Periodic background media scans (Patrol Read) to find and correct media defects to avoid losing data redundancy (hot spare drives only)
- Use of a Commissioned Hotspare to rebuild a degraded drive group (Commissioned Hotspare drives only)
- Update of disk data format (DDF) and other metadata when you make changes to RAID configurations (Commissioned Hotspare drives and unconfigured drives)

NOTE: If your controller does not support this option, the **Power Settings** field does not appear.

Follow these steps to change the power-save setting.

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To>>Controller>>Power Settings** from the menu bar.

The Power Settings dialog appears, as shown in the following figure.

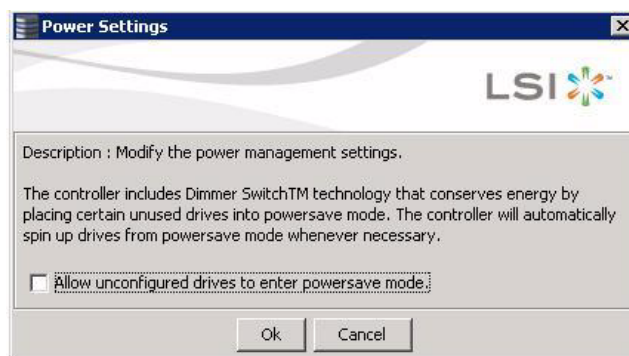


Figure 222: Power Settings Dialog

3. Select the **Allow unconfigured drives to enter Power-Save mode** check box and click **Ok**.

The second Power Settings window appears, as shown in the following figure.

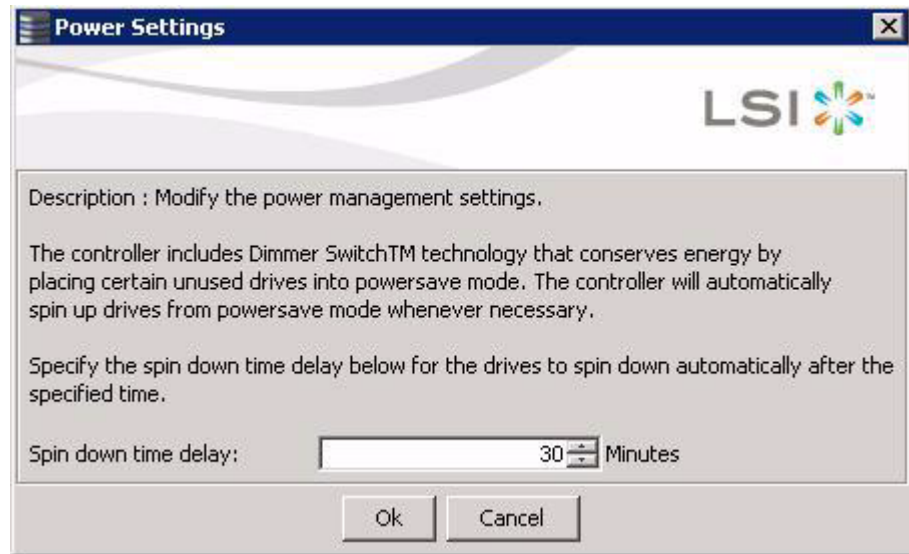


Figure 223: Power Settings Dialog – Spin Down Time Delay Setting

4. Enter the time delay in minutes before the unconfigured drives spin down automatically.

After the specified time, the drives spin down automatically.

5. Click **Ok**.

Your power settings are saved. In the **Physical** tab of the main menu window, the nodes for the unconfigured good drives that are spun down appear with - Powersave after their status.

8.5.1 Enhanced Dimmer Switch Power Settings

This section describes how to change the power-save settings using the Dimmer Switch Enhancement (using the Power-Save mode).

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To>>Controller>>Manage Power Settings** from the menu bar.

The Manage Power Save Settings window appears, as displayed in the following figure.

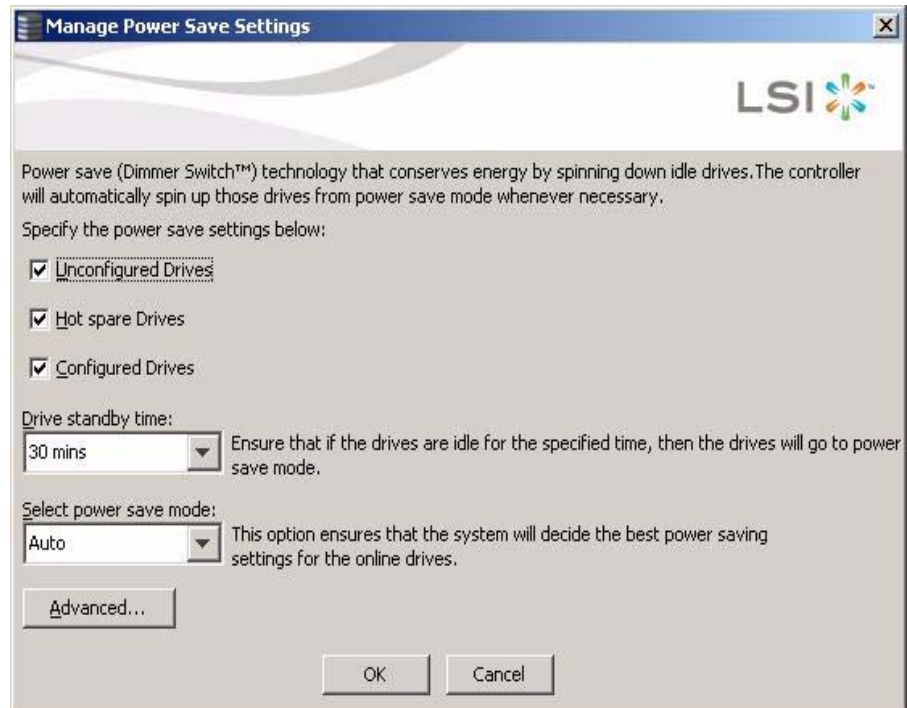


Figure 224: Manage Power Save Settings Window

3. Select the **Unconfigured Drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
4. Select the **Hot spare Drives** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the **Configured Drives** check box to let the controller enable the Configured drives to enter the Power-Save mode.
6. Select the drive standby time (**Alt+D**) using the drop-down list from the **Drive standby time** field.

NOTE: The **Drive Standby time** drop-down list is enabled only if any of the check boxes above it are checked. The drive standby time can be 30 minutes, 1 hour, 1.30 hours, or 2 hours through 24 hours.

7. Select the Power-Save mode using the **Select Power- Save mode** drop down list. The mode can be **Auto**, **Max** or **Max without cache**.

NOTE: The **Select Power-Save mode** drop-down list is enabled only if the **Configured drives** check box is selected. The **Max without cache** mode option depends on the firmware settings.

8. Click **OK**.

The Power-Save settings are saved. After you click **OK**, the following dialog box appears.




Figure 225: Power Save Settings Saved - Confirmation Dialog

If you do not specify the Power-Save settings in the Manage Power Save Settings window as shown in [Figure 224](#), the following dialog appears.




Figure 226: Power-Save Settings - Confirmation Dialog

8.5.2 Power Save Settings - Advanced

You can schedule the drive active time by selecting the **Start time** check box and **End time** check box in the **Power Save Settings - Advanced** window.

Perform the following steps to schedule the drive active time.

1. Click the **Advanced** button in the Manage Power Save Settings window as shown in [Figure 224](#).
2. The Power Save Settings - Advanced window is displayed, as shown in the following figure.



Figure 227: Power Save Settings - Advanced

3. Select the start time and end time using the drop down list from the **Schedule drive active time** field.
4. Click **OK**.

The drive active time for the configured drives is scheduled.

NOTE: Select the **Do not schedule drive active time** check box if you do not want to schedule the drive active time.

8.5.3 Automatically Spin Up Drives

The Dimmer Switch technology also allows the controller to automatically spin up the drives that are in Power-Save mode.

Perform the following steps to access the **Manage Power Save Settings** window:

1. Right-click **Drive group** >> **Manage Power Settings**.

The **Manage Power Save Settings** window appears, as shown in the following figure.



Figure 228: Manage Power Save Settings Dialog

2. Select the Power-Save mode from the drop-down list.

The values can be **Max**, **Max without cache**, **Auto**, **None**, and **controller defined** in the same order.

NOTE: The **Controller defined** option enables the system to inherit the controller Power-Save mode for online drives.

3. Click **OK**.

The Power-Save mode is saved.

8.5.4 Power-Save Mode

The Power-Save mode can be set during creation of the virtual drive by using the **Select power save mode** field in the **Create Drive Group - Drive Group Settings** window as shown in the following figure.

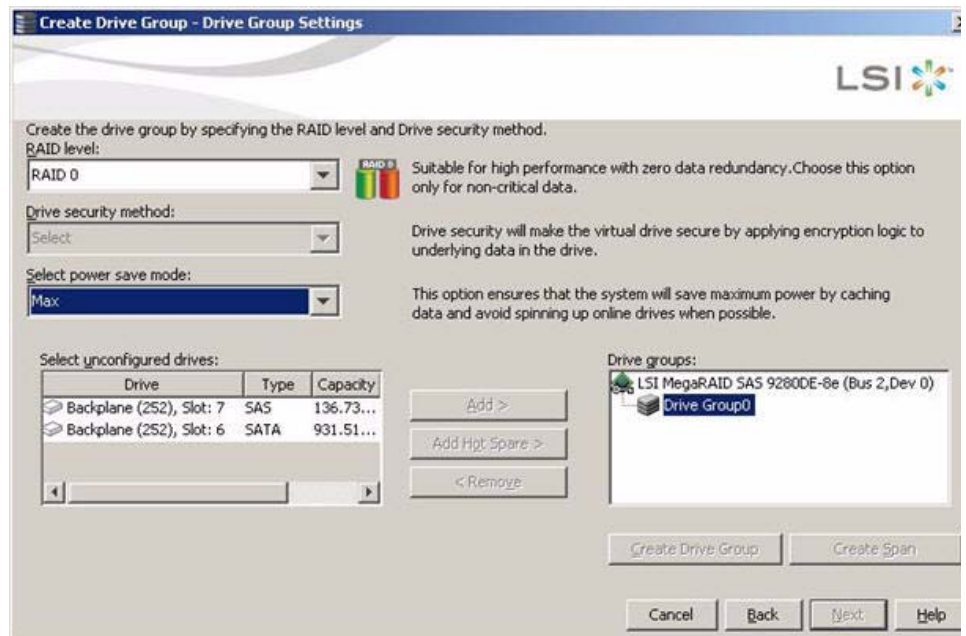


Figure 229: Create Drive Group - Drive Group Settings (Automatic Spin Up)

8.5.5 Power Save Mode - SSD Drives

If you select the **Max** and **Max without cache** options in the **Select power save mode** field in Figure 229, and select one or more SSD drives, and click **Create Drive Group**, the following confirmation dialog appears.



Figure 230: Power Save Mode - Confirmation Dialog

8.6 Changing Virtual Drive Properties

You can change the read policy, write policy, and other virtual drive properties at any time after a virtual drive is created.

CAUTION: Do not enable drive caching on a mirrored drive group (RAID 1 or RAID 1E). If you do, data can be corrupted or lost in the event of a sudden power loss. A warning appears if you try to enable drive caching for a mirrored drive group

NOTE: For virtual drives with SAS drives only, set the drive write cache policy set to *Disabled*, by default. For virtual drives with SATA drives only, set the drive write cache policy to *Enabled*, by default.

To change the virtual drive properties, perform the following steps:

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To>>Virtual Drive>>Set Virtual Drive Properties** from the menu bar, as shown in the following figure.

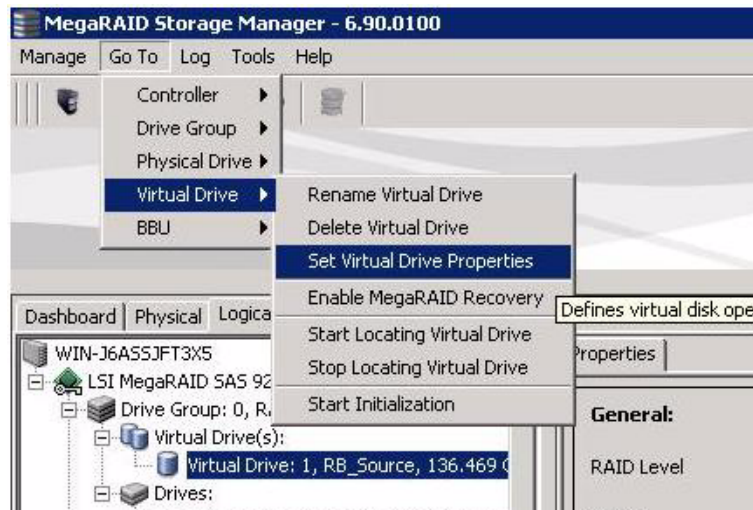


Figure 231: Set Virtual Drive Properties Menu option

The Set Virtual Drive Properties dialog displays, as shown in the following figure.

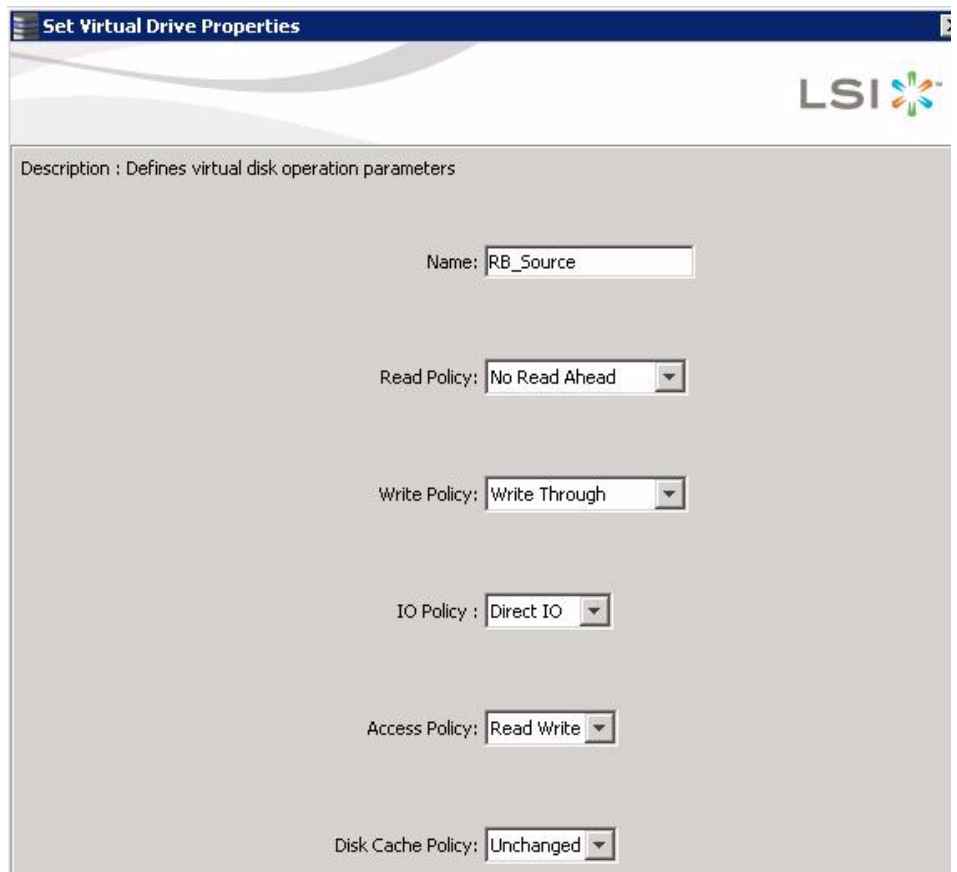


Figure 232: Set Virtual Drive Properties Dialog

3. Change the virtual drive properties as needed.

For information about these properties, see [Section 8.1.1, *Selecting Virtual Drive Settings*](#).

4. Click **Ok** to accept the changes.

The virtual drive settings are updated.

8.7 Changing a Virtual Drive Configuration

You can use the Modify Drive Group wizard in the MegaRAID Storage Manager software to change the configuration of a virtual drive by adding drives to the virtual drive, removing drives from it, or changing its RAID level.

CAUTION: Be sure to back up the data on the virtual drive before you change its configuration.

NOTE: You cannot change the configuration of a RAID 10, RAID 50, or RAID 60 virtual drive. You cannot change a RAID 0, RAID 1, RAID 5, or RAID 6 configuration if two or more virtual drives are defined on a single drive group. (The Logical tab shows which drive groups and drives are used by each virtual drive.)

8.7.1 Accessing the Modify Drive Group Wizard

NOTE: The Modify Drive Group wizard was previously known as the Reconstruction wizard.

Perform the following steps to access the Modify Drive Group wizard options:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu window.
2. Select a drive group in the left panel of the window.
3. Select **Go To>> Drive Group>>Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group wizard.

The following warning appears about rebooting virtual drives containing boot partitions that are undergoing RAID level migration or capacity expansion operations. Back up your data before you proceed.




Figure 233: Reboot Warning Message

4. Select the **Confirm** check box and click **Yes**.

A warning to back up your data appears, as shown in the following figure.




Figure 234: Warning to Back Up Data Message

5. Select the **Confirm** check box and click **Yes**.

The Modify Drive Group wizard window appears, as shown in the following figure.

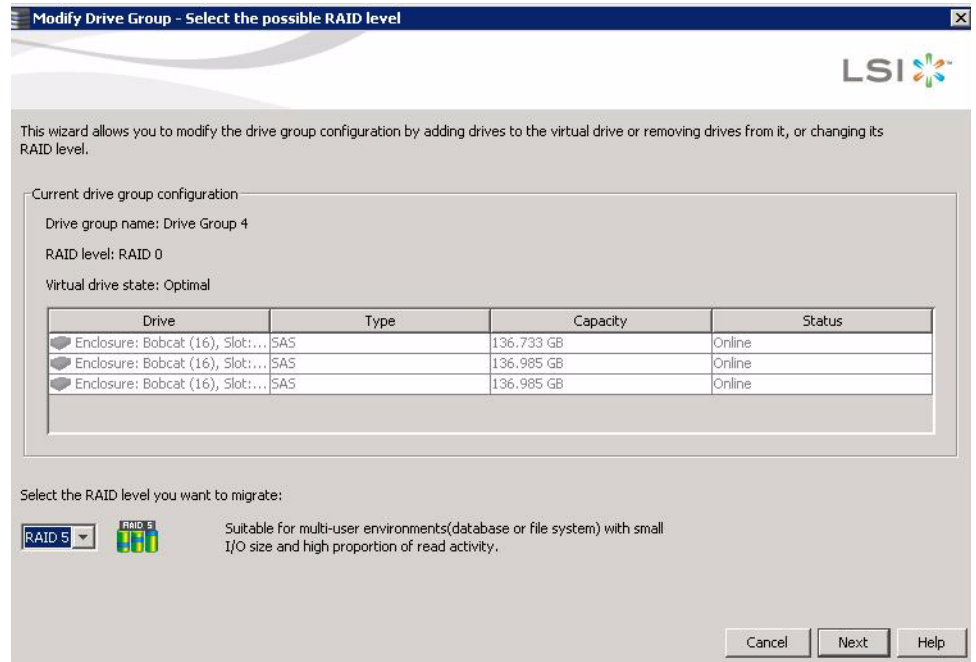


Figure 235: Modify Drive Group Wizard Window

The following sections explain the Modify Drive Group wizard options.

8.7.2 Adding a Drive or Drives to a Configuration

CAUTION: Be sure to back up the data on the virtual drive before you add a drive to it.

Follow these steps to add a drive or drives to a configuration with the Modify Drive Group wizard.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To>> Drive Group>>Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group wizard.

The Modify Drive Group wizard window appears.

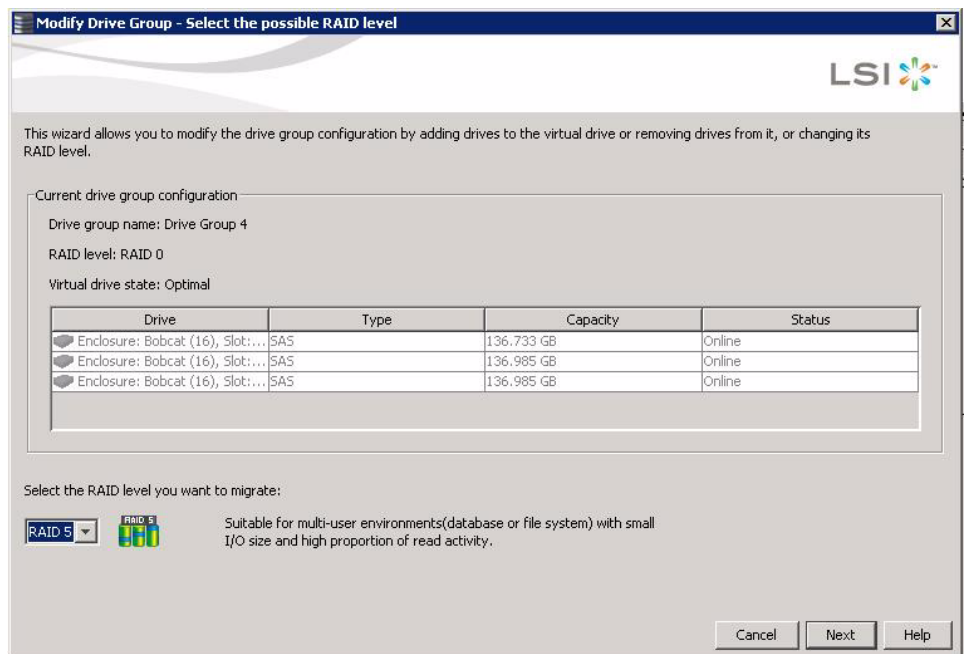


Figure 236: Modify Drive Group Wizard Window

4. Select the RAID level to which you want to change ("migrate") the drive group, and click **Next**.

The following window appears. It lists the drives you can add, and it states whether you have to add a minimum number of drives to change the RAID level from the current level to the new RAID level.

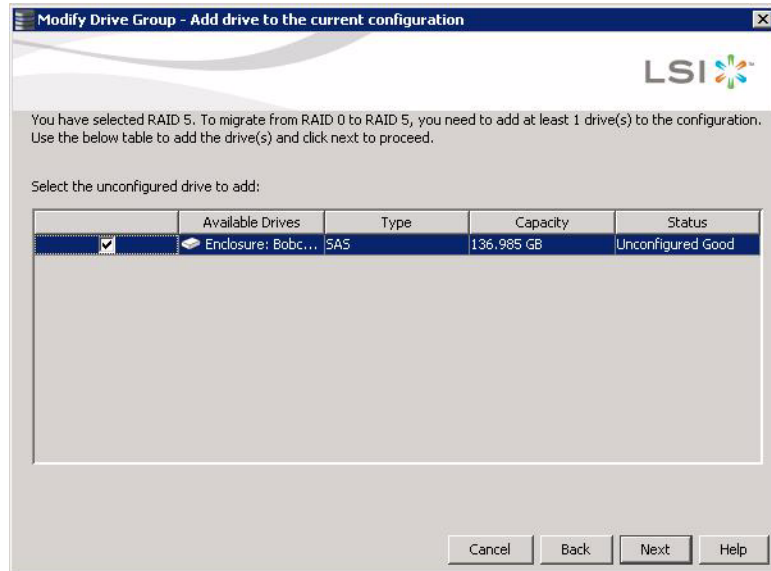


Figure 237: Modify Drive Group – Add Drives to the Current Configuration Window

5. Click the check box next to any unconfigured drives that you want to add, and then click **Next**.

NOTE: The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The Modify Drive Group – Summary window appears. This window shows the current settings and what the settings will be after the drives are added.

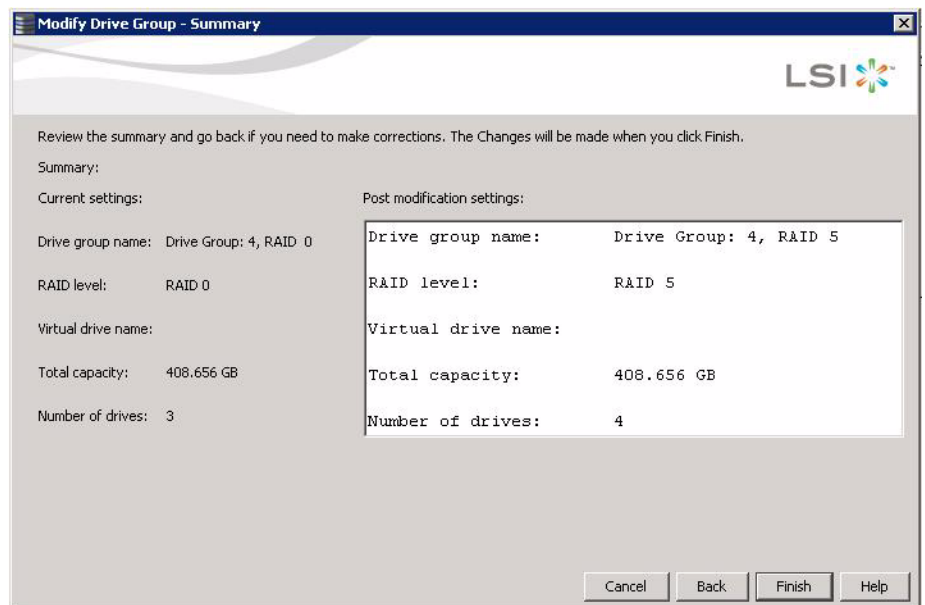


Figure 238: Modify Drive Group – Summary Window

6. Review the configuration information.

You can click **Back** if you need to change any selections.

7. Click **Finish** to accept the changes.

A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

8. Click **Yes** to accept and complete the addition of the drives to the drive group.

8.7.3 Removing a Drive from a Configuration

CAUTION: Be sure to back up the data on the virtual drive before you remove a drive from it.

Follow these steps to remove a drive from a RAID 1, RAID 5, or RAID 6 configuration.

NOTE: This option is not available for RAID 0 configurations.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Click a drive icon in the left panel of the window.
3. Either select **Go To>> Physical Drive>>Make Drive Offline** on the menu bar, or right-click the drive, and select **Make Drive Offline** from the menu.

A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

4. Click **Yes** to accept and complete the removal of the drive from the drive group.

8.7.4 Replacing a Drive

CAUTION: Be sure to back up the data on the virtual drive before you replace a drive.

Follow these steps to add a replacement drive and copy the data from the drive that was removed to the replacement drive.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive in the left panel of the window.
3. Either select **Go To>> Physical Drive>>Replace Physical Drive** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group wizard.

The dialog with the replacement drive appears, as shown in the following figure.

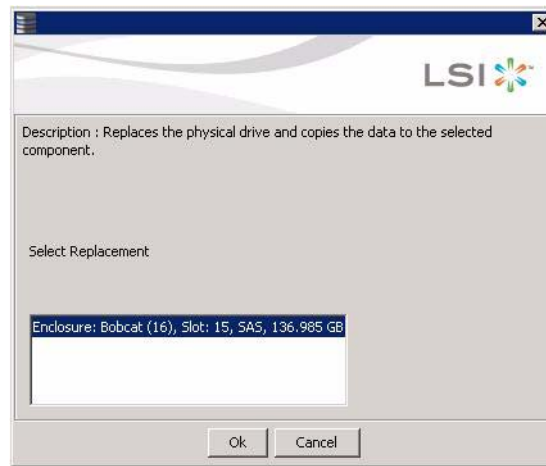


Figure 239: Drive Replacement Window

4. Select a replacement drive.

A confirmation message appears.

5. Click **Yes**.

This step replaces a drive and copies the data to the selected component.

8.7.5 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 make this change.

When you migrate a virtual drive to another RAID level, you can keep the same number of drives, or you can add drives. In some cases, you have to add a certain number of drives to migrate the virtual drive from one RAID level to another. The window indicates the minimum number of drives you are required to add.

CAUTION: Be sure to back up the data on the virtual drive before you change the RAID level.

Follow these steps to change the RAID level of the virtual drive with the Modify Drive Group wizard:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To>> Drive Group>>Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group wizard.

The Modify Drive Group wizard appears.

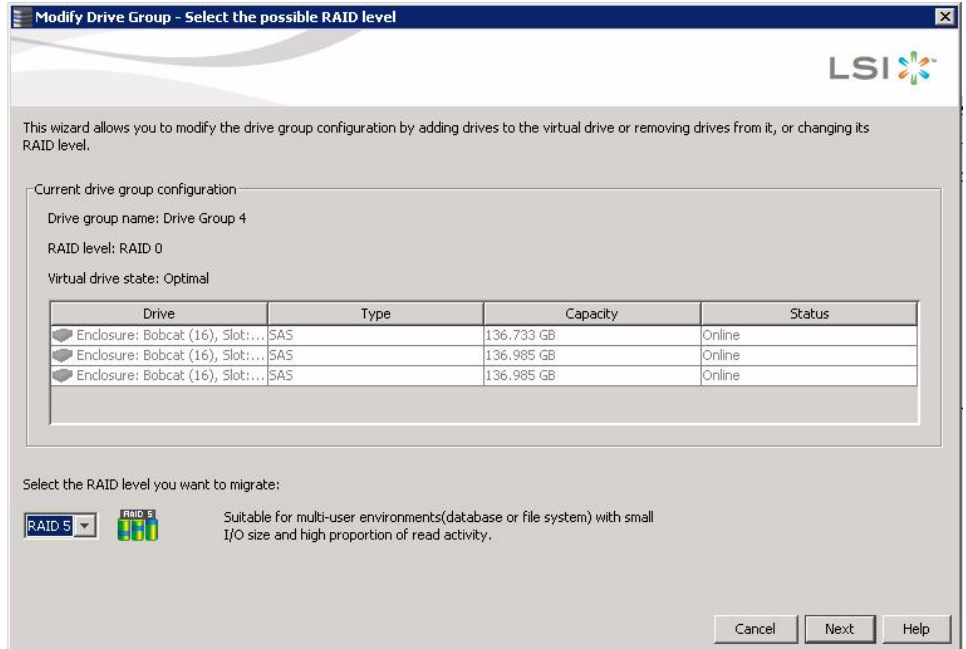


Figure 240: Modify Drive Group Wizard Dialog

- On the Modify Drive Group Wizard dialog, select the RAID level to which you want to change ("migrate") the drive group to, and click **Next**.

The following dialog appears. The dialog states the number of drives that you have to add to change the RAID level from the current level to a new RAID level that requires more drives.

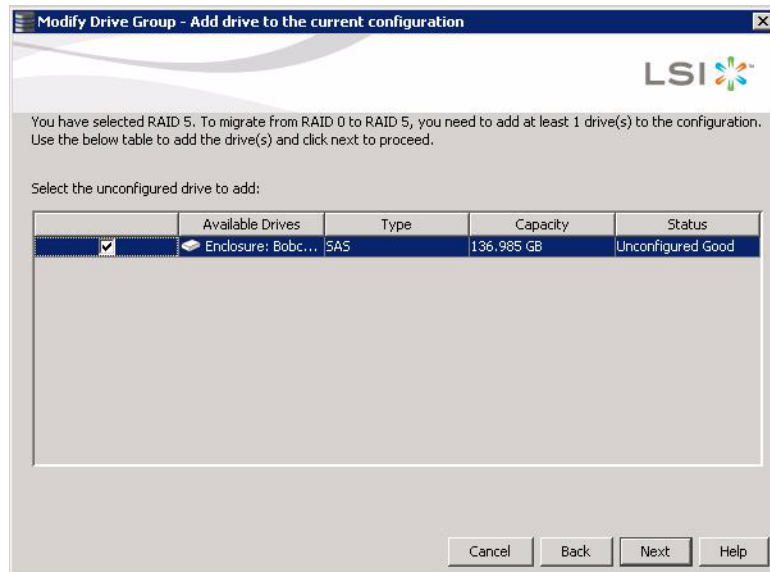


Figure 241: Modify Drive Group -Add Drives to the Current Configuration Screen

- Select the unconfigured drive or drives to add, and click **Next**.

NOTE: The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The Modify Drive Group – Summary window appears. This window shows the current settings and what the settings will be after the drives are added.

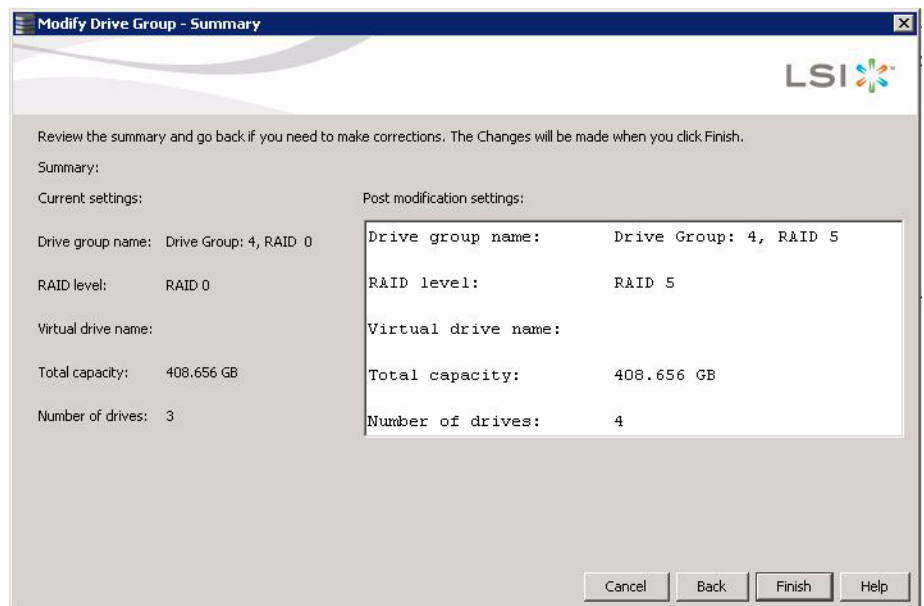


Figure 242: Modify Drive Group – Summary Screen

- Review the configuration information.

You can click **Back** if you need to change any selections.

- Click **Finish** to accept the changes.

A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

- Click **Yes** to accept and complete the migration to the new RAID level.

The operation begins on the virtual disk. To monitor the progress of the RAID level change, select **Manage >> Show Progress** in the menu bar.

8.7.6 New Drives Attached to a MegaRAID Controller

When you insert a new drive on a MegaRAID system, if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD for MegaRAID entry-level controllers, such as the SAS 9240-4i/8i. If the drive does contain valid DDF metadata, its drive state is Unconfigured Good.

A new drive in JBOD drive state is exposed to the host operating system as a stand-alone drive. Drives in JBOD drive state are not part of the RAID configuration because they do not have valid DDF records. The operating system can install and run anything on JBOD drives.

Automatic rebuilds always occur when the drive slot status changes, for example, when you insert a drive or remove a drive, so that a Commissioned Hotspare can be used. However, a new drive in JBOD drive state (without a valid DDF record), does not perform an automatic rebuild.

To start an automatic rebuild on the new JBOD drive, you have to change the drive state from JBOD to Unconfigured Good. (Rebuilds start only on Unconfigured Good drives.) After you set the drive state to Unconfigured Good, the drive state information always remains on the drive, and you can use the drive for configuration.

See [Section 4.14.3.3, Troubleshooting Information](#) for more information about DDF and metadata. See [Section 10.5, Making a Drive Offline or Missing](#) for the procedure to change a drive to the Unconfigured Good drive state.

8.8 Deleting a Virtual Drive

CAUTION: Make sure to back up the data that is on the virtual drive before you delete it. Make sure that the operating system is not installed on this virtual drive.

You can delete virtual drives to rearrange the storage space. To delete a virtual drive, follow these steps.

1. Back up all user data that is on the virtual drive you want to delete.
2. On the MegaRAID Storage Manager window, select the **Logical** tab, and click the icon of the virtual drive you want to delete.
3. Select **Go To>>Virtual Drive>>Delete Virtual Drive**.
4. When the warning messages appear, click **Yes** to confirm that you want to delete the virtual drive.

NOTE: You are asked twice if you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

Chapter 9

Monitoring Controllers and Its Attached Devices

This chapter explains how to use the MegaRAID Storage Manager software to monitor the status of drives, virtual drives, and other storage devices.

The MegaRAID Storage Manager software enables you to monitor the activity of all the controllers present in the system and the devices attached to them.

When you perform an operation on devices (such as the creation of a new virtual drive) or when devices automatically go from an optimal state to a different state (such as a created virtual drive goes to a degraded state or a Battery Backup Unit goes bad), the MegaRAID Storage Manager software gets those events from the controller and gives a notification to you, using different alert delivery methods

9.1 Alert Delivery Methods

Based on the severity level (Information, Warning, Critical and Fatal), the default alert delivery methods change. By default, each severity level has one or more alert delivery methods configured for it, as shown in [Table 148](#). To modify these alert delivery methods, see [Configuring Alert Notifications](#). The different alert delivery methods are as follows:

- Vivaldi Log/MegaRAID Storage Manager Log
- System Log
- Pop-up Notification
- E-mail Notification

Table 148: Severity Level and Default Alert Delivery Methods

Severity Level	Default Alert Delivery Method	Meaning
Information	Vivaldi log/MegaRAID Storage Manager log and System log	Informational message. No user action is necessary.
Warning	Vivaldi log/MegaRAID Storage Manager log and System log	Some component might be close to a failure point.
Critical	Vivaldi log/MegaRAID Storage Manager log, System log, and Popup Notification	A component has failed, but the system has not lost data.
Fatal	Vivaldi log/MegaRAID Storage Manager log, System log, Popup Notification, and E-mail Notification	A component has failed, and data loss has occurred or will occur.

9.1.1 Vivaldi Log / MegaRAID Storage Manager Log

By default, all the severity events appear in the Vivaldi log/MegaRAID Storage Manager log and are displayed at the bottom of the MegaRAID Storage Manager main menu window. Each message that appears in this log has a severity level that indicates the importance of the event (severity), a date and timestamp (when it occurred), and a brief description, as show in the following figure.

ID	Error Level	Date / Time	Description
247	[Information, 0]	2011-03-17, 19:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,14:00:48 1465794 Seconds
246	[Information, 0]	2011-03-17, 18:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,13:00:48 1462194 Seconds
245	[Information, 0]	2011-03-17, 12:32:06	Successful log on to the server User: root, Client: 135.24.228.164, Access Mode: Full, Client Time: 2011-03-17,13:51:46
244	[Information, 0]	2011-03-17, 12:11:25	Successful log off from the server User: root, Client: 135.24.228.164, Client Time: 2011-03-17,13:31:09
243	[Information, 0]	2011-03-17, 17:41:17	Controller ID: 0 Configuration cleared
242	[Information, 0]	2011-03-17, 17:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,12:00:48 1458594 Seconds
241	[Information, 0]	2011-03-17, 16:49:43	Controller ID: 0 Configuration cleared
240	[Information, 0]	2011-03-17, 11:16:24	Successful log on to the server User: root, Client: 135.24.228.164, Access Mode: Full, Client Time: 2011-03-17,12:36:04
239	[Information, 0]	2011-03-17, 11:04:11	Successful log out from the server User: root, Client: 135.24.228.117, Client Time: 2011-03-17,11:04:11
238	[Information, 0]	2011-03-17, 16:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,11:00:48 1454994 Seconds
237	[Information, 0]	2011-03-17, 16:17:06	Controller ID: 0 Rebuild rate changed: 36 %
236	[Information, 0]	2011-03-17, 16:11:25	Controller ID: 0 Configuration cleared
235	[Information, 0]	2011-03-17, 10:40:24	Successful log on to the server User: root, Client: 135.24.228.117, Access Mode: Full, Client Time: 2011-03-17,10:40:24
234	[Information, 0]	2011-03-17, 15:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,10:00:48 1451394 Seconds
233	[Information, 0]	2011-03-17, 14:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,09:00:48 1447794 Seconds
232	[Information, 0]	2011-03-17, 13:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,08:00:48 1444194 Seconds
231	[Information, 0]	2011-03-17, 12:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,07:00:48 1440594 Seconds
230	[Information, 0]	2011-03-17, 11:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,06:00:48 1436994 Seconds
229	[Information, 0]	2011-03-17, 10:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,05:00:48 1433394 Seconds
228	[Information, 0]	2011-03-17, 09:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,04:00:48 1429794 Seconds
227	[Information, 0]	2011-03-17, 08:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,03:00:48 1426194 Seconds
226	[Information, 0]	2011-03-17, 07:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,02:00:48 1422594 Seconds
225	[Information, 0]	2011-03-17, 06:30:48	Controller ID: 0 Time established since power on: Time: 2011-03-17,01:00:48 1418994 Seconds

Figure 243: Vivaldi Log

You can double click on an event to display the same information in a separate window. For a list of all events, see [Events and Messages](#). The status bar at the bottom of the screen indicates whether the log is a MegaRAID Storage Manager server log or a locally stored log file.

When a Vivaldi log/MegaRAID Storage Manager log appears, the Log menu has the following options:

- **Save Log:** Saves the current log to a .log file.
- **Save Log Text:** Saves the current log in .txt format.
- **Load:** Enables you to load a local .log file in the bottom of the MegaRAID Storage Manager main menu window. If you select the **Load** menu, you will not be able to view the current log.
- **Rollback to Current Log:** This menu appears if we have loaded the logs from a local .log file. Once you select this menu, you can view the current log.
- **Clear Log:** Clears the current log information, if you have full access (versus view-only access). You have the option to save the log first.

9.1.2 System Log

By default, all the severity events are logged in the local syslog. Based on the operating system you are using, the system log is logged in the following syslog locations:

- In Windows, the system log is logged in Event Viewer >> Application.
- In Linux, the system log is logged in /var/log/messages.
- In Solaris, the system log is logged in /var/adm/messages.

9.1.3 Pop-up Notification

By default, fatal and critical events are displaying in a pop-up notification. Pop-up notification is started automatically when you are login in to the operating system. Through this feature, you can view multiple events in a single pop-up window as shown in following figure.

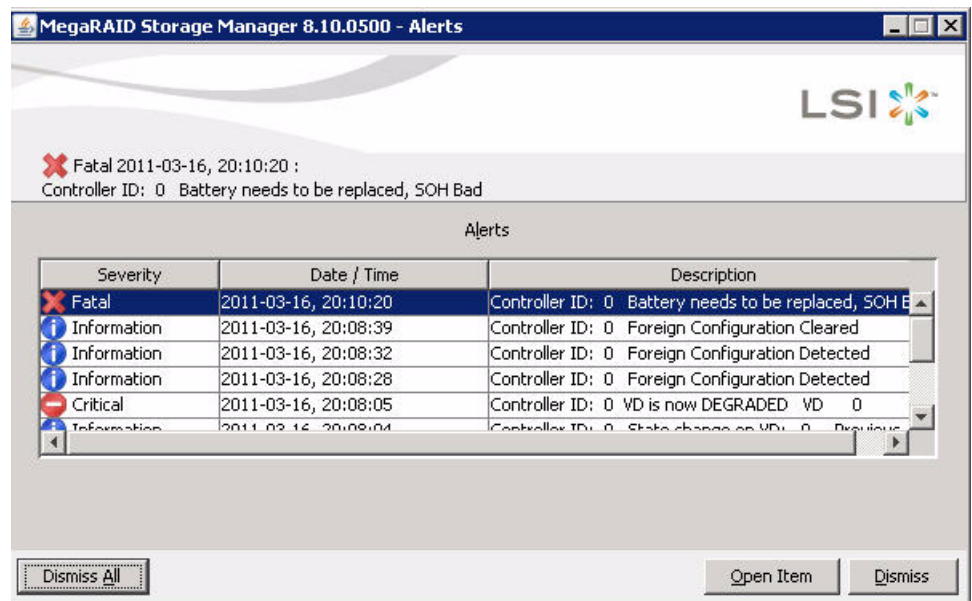


Figure 244: Pop-Up Notification

9.1.4 E-mail Notification

By default, fatal events are displayed as e-mail notifications. Based on your configuration, the e-mail notifications are delivered to you as shown in the following figure.

In the e-mail notification, besides the event's description, the email also contains system information and the controller's image details. Using this additional information, you can find out the system and the controller on which the fatal error occurred.

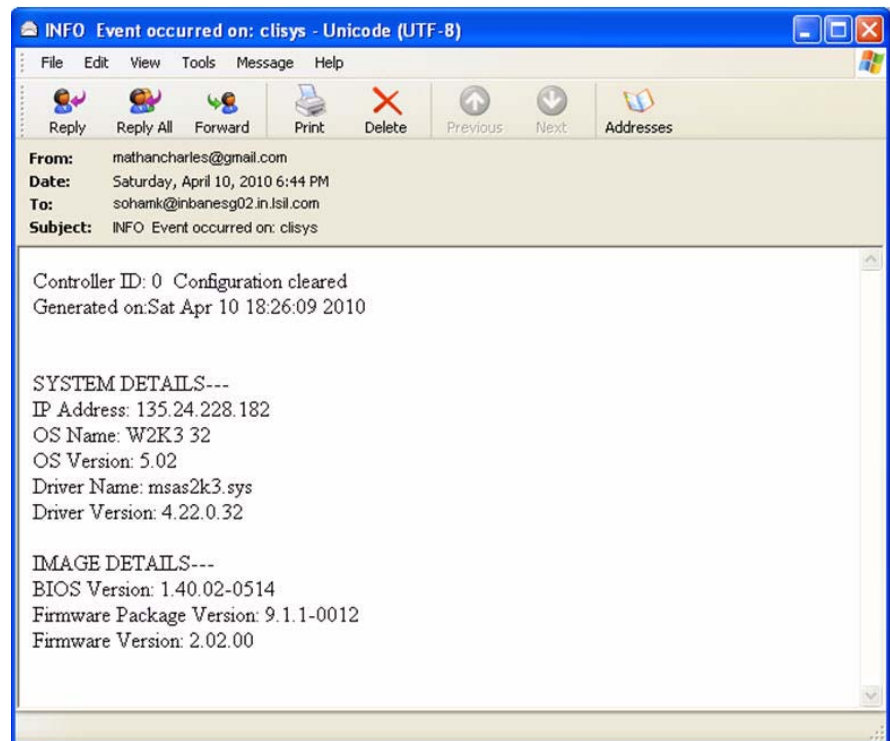


Figure 245: E Mail Notification

9.2 Configuring Alert Notifications

The Alert Notification Configuration feature allows you to control and configure the alerts that the MegaRAID Storage Manager software sends when various system events occur.

Select **Tools->Configure Alerts** on the main menu screen, as shown in the following figure.



Figure 246: Alert Notification Configuration Menu

The Configure Alerts window appears, as shown in the following figure. The window contains three tabs: **Alert Settings**, **Mail Server**, and **Email**.

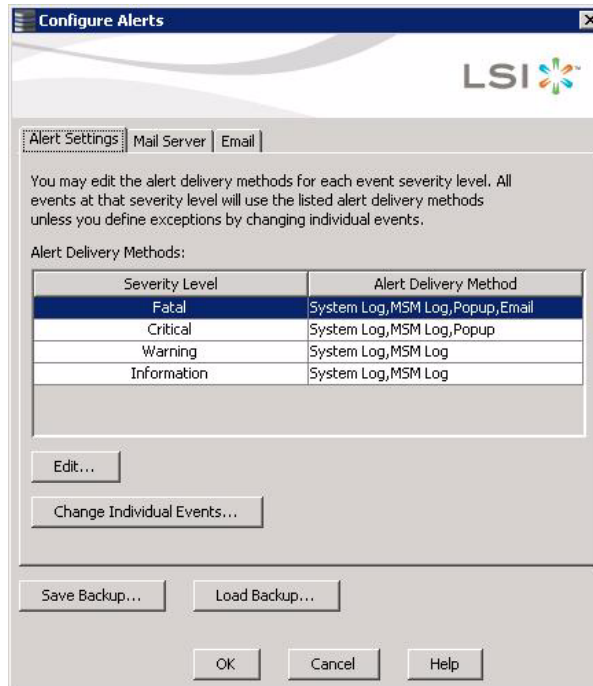


Figure 247: Configure Alerts

You can select the **Alert Settings** tab to perform the following actions:

- Edit the alert delivery method for different severity levels.
- Change the method of delivery for each individual event.
- Change the severity level of each individual event.
- Save an .xml backup file of the entire alert configuration.
- Load all the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

NOTE: When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender email address.
- Enter the SMTP server name or the IP address.
- Enter the SMTP server authentication related information (user name and password).

NOTE: These fields are optional and are filled only when the SMTP server requires authentication.

- Save an .xml backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

NOTE: When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Email** tab to perform the following actions:

- Add new email addresses for recipients of alert notifications.
- Send test messages to the recipient email addresses.
- Remove email addresses of recipients of alert notifications.
- Save an .xml backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

NOTE: When you load a saved backup file, all unsaved changes made in the current session will be lost.

9.3 Editing Alert Delivery Methods

You can edit the default alert delivery methods, such as pop-up, email, system log, or the Vivaldi Log / MegaRAID Storage Manager log to different severity level (Information, Warning, Critical and Fatal).

Perform the following steps to edit the alert delivery methods:

1. On the Configure Alerts window, click the **Alerts Setting** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**. The Edit dialog appears, as shown in the following figure.



Figure 248: Edit Dialog

4. Select the desired alert delivery methods for alert notifications at the event severity level.
5. Click **OK** to set the delivery methods used for the severity level that you selected.

9.4 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.

1. On the Configure Alerts window, click the **Alerts Setting** tab.

The **Alerts Setting** portion of the window appears.

2. Click **Change Individual Events**.

The **Change Individual Events** dialog appears, as shown in the following figure. The dialog shows the events by their ID number, description, and the severity level.

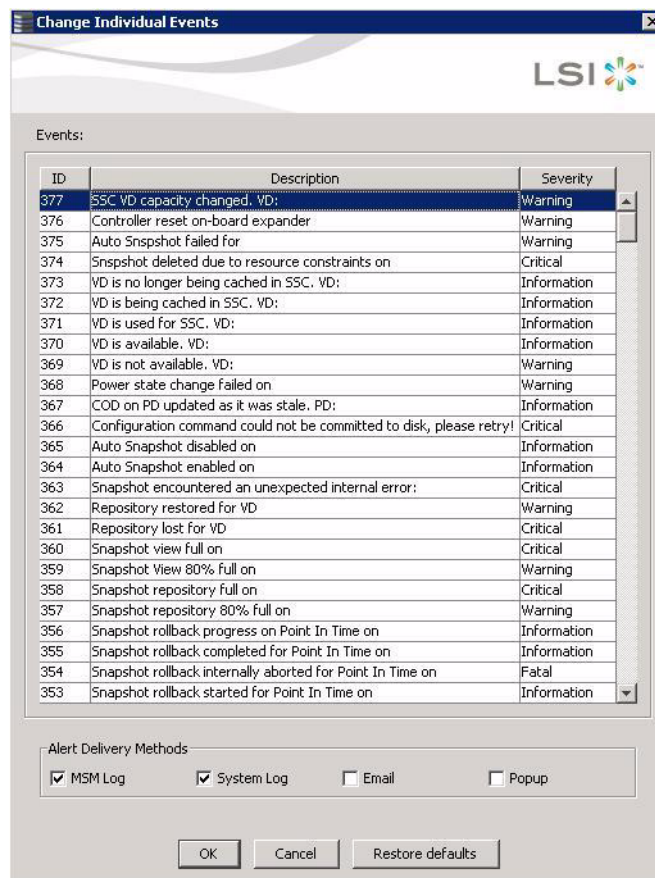


Figure 249: Change Individual Events

3. Click an event in the list to select it.

The current alert delivery methods appear for the selected event in the **Alert Delivery Methods** frame.

4. Select the desired alert delivery methods for the event.
5. Click **OK** to return to the **Configure Alerts** window.
6. You may click **Cancel** to discard your current changes and to go back to the **Configure Alerts** window.
7. In the **Configure Alerts** window, click **OK**. This saves all the changes made to the event.

NOTE: You can click **Restore Defaults** to revert back to the default alert delivery method and the default severity level of an individual event. For more information, see [Rollback to Default Individual Event Configuration](#).

9.5 Changing the Severity Level for Individual Events

To change the event severity level for a specific event, perform the following steps:

NOTE: See [Table 148](#) for details about the severity levels.

1. On the Configure Alerts window, click the **Alerts Setting** tab.
The **Alerts Setting** portion of the window appears.
2. Click **Change Individual Events**. The **Change Individual Events** dialog appears.
The dialog shows the events by their ID number, description, and severity level.
3. Click an event in the list to select it. The current severity appears in the Severity cell for the selected event.
4. Click the **Severity** cell for the event. The Event Severity drop-down menu appears for that event, as shown in the following figure.

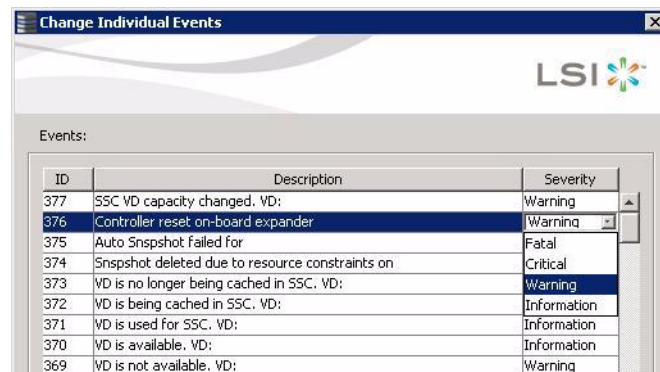


Figure 250: Change Individual Events Severity Level Menu

5. Select a different severity level for the event from the menu.
6. Click **OK** to return to the **Configure Alerts** window.
7. You may click Cancel to discard your current changes and to go back to the **Configure Alerts** window.
8. In the **Configure Alerts** window, click **OK** to save all the changes made to the events.

9.6 Rollback to Default Individual Event Configuration

To revert back to the default alert delivery method and the default severity level of an individual event, perform the following steps:

1. On the **Configure Alerts** window, click the Alerts Setting tab.
The Alerts Setting portion of the window appears.
2. Click **Change Individual Events**.
The Change Individual Events dialog appears, as shown in [Figure 249](#). The dialog shows the events by their ID number, description, and the severity level.
3. Click **Restore Defaults**.
The Change Individual Events dialog appears with the default alert delivery method and the default severity level of all individual events.
4. Click OK to return to the **Configure Alerts** window.
5. In the **Configure Alerts** window, click OK to save all the changes made to the events.

9.7 Entering or Editing the Sender Email Address and SMTP Server

You can use the **Configure Alerts** window to enter or edit the sender e-mail address and the SMTP server.

1. On the **Configure Alerts** window, click the **Mail Server** tab.

The **Mail Server** options appear, as shown in the following figure.



Figure 251: Mail Server Options

2. Enter a sender's email address in the **Sender email address** field or edit the existing sender email address.
3. Enter your SMTP server name/IP Address in the **SMTP Server** field or edit the existing details.
4. Click **OK**.

9.8 Authenticating the SMTP Server

The MegaRAID Storage Manager software supports a SMTP authentication mechanism called *Login*. This feature provides an extra level of security, while sending an email from the MegaRAID Storage Manager server.

To enter or modify the SMTP server authentication information, perform the following steps:

1. On the **Configure Alerts** window, click the **Mail Server** tab.
The **Mail Server** options appear, as shown in [Figure 251](#).
2. If on your SMTP server, the authentication mechanism is enabled and if you want to enable this feature on the MegaRAID Storage Manager software, then you need to select the **This Server requires authentication** check box and enter the authentication details in the corresponding fields (**User name** and **Password**).

If you do not want to enable this feature on the MegaRAID Storage Manager software or if you know that your SMTP server does not support the *Login* mechanism, then de-select the **This Server requires authentication** check box.

NOTE: The **This Server requires authentication** check box is selected by default.

3. Enter a user name in the **User name** field. (Optional - if **This Server requires authentication** check box is selected).
4. Enter the password in the **Password** field. (Optional - if **This Server requires authentication** check box is selected).
5. Click **OK**.

9.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab in the Configure Alerts window shows the email addresses of the recipients of the alert notifications. The MegaRAID Storage Manager software sends alert notifications to those email addresses. Use the Configure Alerts window to add or remove email addresses of recipients and to send test messages to recipients that you add.

To add email addresses of recipients of the alert notifications, perform the following steps:

1. Click the **E-mail** tab in the Configure Alerts window.

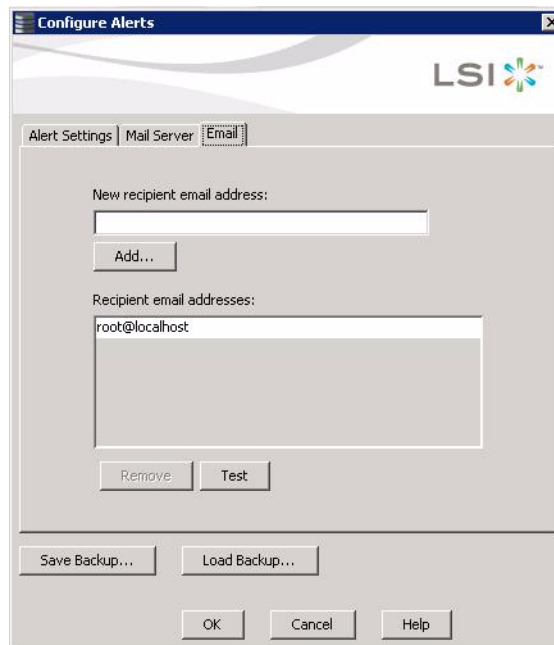


Figure 252: Adding Email Settings

2. Enter the email address you want to add in the **New recipient email address** field.
3. Click **Add**. The new email address appears in the **Recipient email addresses** field.

9.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab in the Configure Alerts window to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **Email** tab on the Configure Alerts window.

The **Email** section of the window appears, as shown in the above figure.

2. Click an email address in the **Recipient email addresses** field.
3. Click **Test**.
4. Confirm whether the test message was sent to the email address.

A pop-up message indicates if the test message sent to the email address was successful. If the MegaRAID Storage Manager software cannot send an email message to the email address, an error message appears.

9.11 Removing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab in the Configure Alerts window to remove email addresses of the recipients of alert notifications.

1. Click the **Email** tab on the Configure Alerts window.
The **Email** section of the window appears, as shown in the previous figure.
2. Click an email address in the **Recipient email addresses** field.
The **Remove** button, which was grayed out, is now active.
3. Click **Remove**. The email address is deleted from the list.

9.12 Saving Backup Configurations

You can save an `.xml` backup file of the entire alert configuration. This includes all the settings on the three tabs (Alert Settings, Mail Server, and Email).

1. On the **Configure Alerts** window, click the **Alert Setting** tab, **Mail Server** tab, or the **Email** tab.
2. Click **Save Backup**. The drive directory appears.
3. Enter a filename with an `.xml` extension for the backup configuration (in the format `filename.xml`).
4. Click **Save**. The drive directory disappears.
5. Click **OK**. The backup configuration is saved and the **Configure Alerts** window closes.

9.13 Loading Backup Configurations

You can load all of the values from a previously saved backup into the **Configure Alert** window (all tabs) to edit or save these values as the current alert notification configuration.

NOTE: If you choose to load a backup configuration and the Configure Alerts window currently contains changes that have not yet been saved as the current alert notification configuration, the changes will be lost. You are prompted to confirm your choice.


1. On the **Configure Alerts** window, click the **Alert Setting** tab, **Mail Server** tab, or **Email** tab.
2. Click **Load Backup**.
You are prompted to confirm your choice. The drive directory appears from which you can select a backup configuration to load.
3. Select the backup configuration file (it should be in `.xml` format).
4. Click **Open**. The drive directory disappears.
5. Click **OK**. The backup configuration is saved and the **Configure Alerts** window closes.

9.14 Monitoring Server Events

The MegaRAID Storage Manager software enables you to monitor the activity of MegaRAID Storage Manager users in the network.

When a user logs on/logs off from the application, the event message appears in the log displayed at the bottom of the MegaRAID Storage Manager screen (the Vivaldi log/MegaRAID Storage Manager Log). These event message have a severity level, a date and timestamp (User log on / log off time), and a brief description that contains a user name, client IP address, an access mode (full/view only) and a client system time.

9.15 Monitoring Controllers

When the MegaRAID Storage Manager software is running, you can see the status of all the controllers in the left panel. If a controller is operating normally, the controller icon looks like this: . If a controller has failed, a small red circle appears next to the icon.

To display the complete controller information, click on a controller icon in the left panel of the MegaRAID Storage Manager main menu. The controller properties appear in the right panel as shown in the following figure. Most of the information on this tab is self-explanatory.

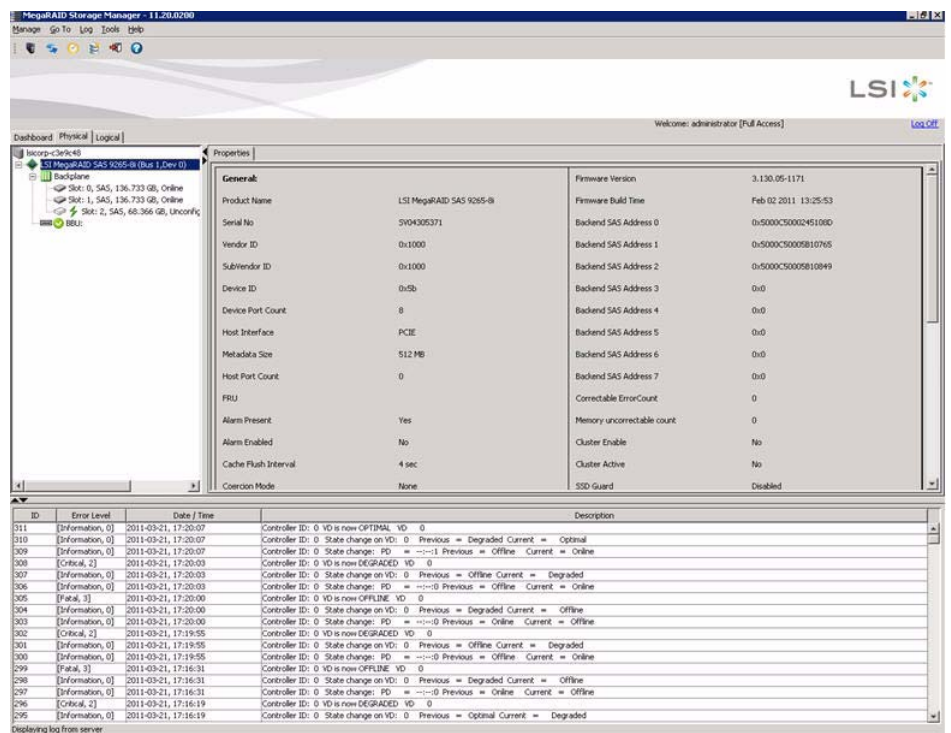


Figure 253: Controller Properties

The Rebuild rate, Patrol read rate, Reconstruction rate, Consistency check rate, and BGI rate (background initialization) are all user selectable. For more information, see [Section 8.4, Changing Adjustable Task Rates](#).

The **BBU Present** field indicates whether a battery backup unit is installed or not.

If the drives are in a disk enclosure, you can identify which drive is represented by a disk icon on the left. To do this, follow these steps:

1. Click the drive icon in the left panel.
2. Select **Go To->Physical Drive->Start Locating Drive** tab in the right panel.

The LED on the drive in the enclosure starts blinking to show its location.

NOTE: LEDs on drives that are global hot spares do not blink.

3. To stop the drive light on the enclosure from blinking, select **Go To->Physical Drive->Stop Locating Drive**.

9.17 Running a Patrol Read

A patrol read periodically verifies all sectors of the drives connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined period and has no other background activities.

You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.

1. Click a controller icon in the left panel.
2. Select **Go To->Controller->Set Patrol Read Properties** or right-click on a controller and select **Set Patrol Read Properties** from the menu.

The Patrol Read - Set properties window appears, as shown in the following figure.

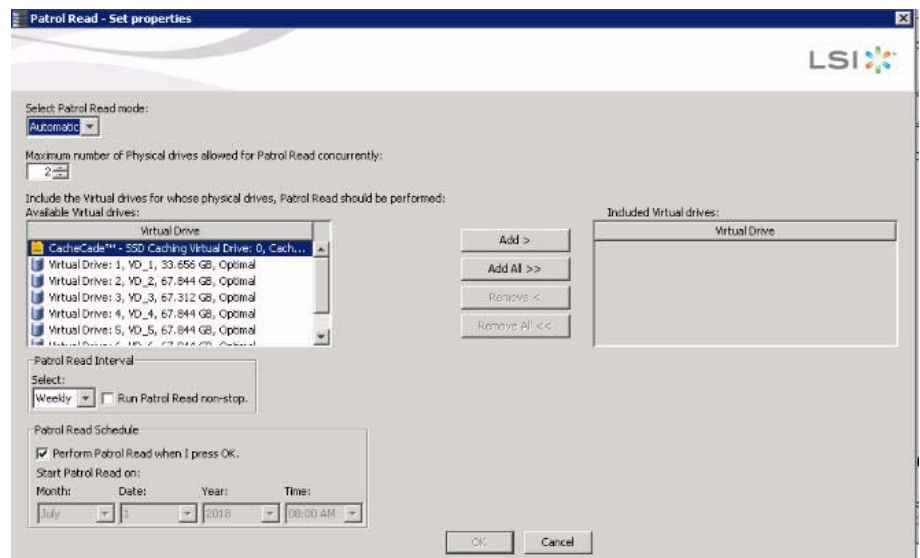


Figure 255: Patrol Read- Set Properties

1. Select an operation mode for patrol read from the following options:
 - **Automatic:** Patrol read runs automatically at the time interval you specify on this window.
 - **Manual:** Patrol read runs only when you manually start it, by selecting **Start Patrol Read** from the controller options window.
 - **Disabled:** Patrol read does not run.
2. (Optional) Specify a maximum count of drives to include in the patrol read.
The count must be a number from 1 to 255.
3. (Optional) Click virtual drives in the list under the heading **Virtual Drives** to include in the patrol read and click **Add >** or click **Add All >>** to include all of the virtual drives.
4. (Optional) Change the frequency at which the patrol read runs.

The default frequency is weekly (168 hours), which is suitable for most configurations. The other options are hourly, daily, and monthly.

NOTE: LSI recommends that you leave the patrol read frequency and other patrol read settings at the default values to achieve the best system performance. If you decide to change the values, record the original default values here so you can restore them later, if necessary:

Patrol Read Frequency: _____

Continuous Patrolling: Enabled/Disabled

Patrol Read Task Rate: _____

5. (Optional) Set Patrol Read to run at a specific time.
The default setting for the patrol read is to start when you click **OK** on this window. To change the default setting so that the patrol read starts at a specific time, follow these steps (otherwise, skip this step and proceed to step 8):
 - a. Deselect the **Perform Patrol Read when I click OK** check box.
 - b. Select the month, year, day, and time to start the patrol read.
6. Click **OK** to enable your patrol read selections.

NOTE: Patrol read does not report on its progress while it is running. The patrol read status is reported only in the event log.

7. Click **Go** to enable these Patrol Read options.

To start a patrol read without changing the patrol read properties, follow these steps:


1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To->Controller->Start Patrol Read** in the menu bar or right-click a controller and select **Start Patrol Read** from the menu.
3. When prompted, click **Yes** to confirm that you want to start a patrol read.

9.17.0.1 Patrol Read Task Rates

You have the option to change the patrol read *task rate*. The task rate determines the amount of system resources that are dedicated to a patrol read when it is running. LSI recommends, however, that you leave the patrol read task rate at its default setting.

If you raise the task rate above the default, the foreground tasks will run more slowly and it may seem that the system is not responding. If you lower the task rate below the default, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. For more information, about the patrol read task rate, see [Changing Adjustable Task Rates](#).

9.18 Monitoring Virtual Drives

When the MegaRAID Storage Manager software is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this: . Color-coded circles appear next to the icon to indicate the following:

- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (for example, if a drive has failed); the data is still safe, but data could be lost if another drive fails.
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

When the Logical tab is selected, the panel on the left shows which drives are used by each virtual drive. The same drive can be used by multiple virtual drives.

To display complete virtual drive information, click the **Logical** tab in the left panel, and click on a virtual drive icon in the left panel. The properties appear in the right panel as shown in the following figure. The RAID level, strip size, and access policy of the virtual drive are set when the virtual drive is configured.

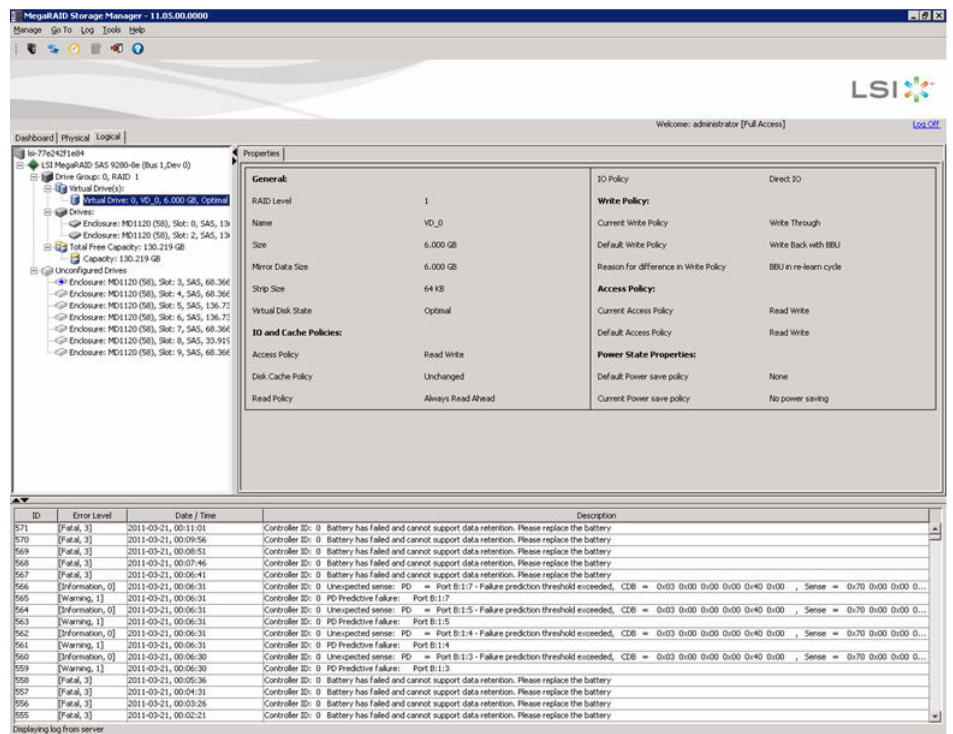


Figure 256: Virtual Drive Properties

You can change the read policy, write policy, and other virtual drive properties. To change these properties, see [Changing Virtual Drive Properties](#).

NOTE: You can change the Read Policy, Write Policy, and other virtual drive properties by selecting the virtual drive icon and then selecting **Go To->Virtual Drive ->Set Virtual Drive Properties** in the menu bar.


If the drives in the virtual drive are in a disk enclosure, you can identify them by making their LEDs blink. To identify the drives, follow these steps:

1. Click the virtual drive icon in the left panel.
2. Either select **Go To->Virtual Drive->Start Locating Virtual Drive**, or right-click a virtual drive and select **Start Locating Virtual Drive** from the menu.

The LEDs on the drives in the virtual drive start blinking (except for the hot spare drives).

3. To stop the LEDs from blinking, select **Go To->Virtual Drive->Stop Locating Virtual Drive** or right-click a virtual drive and select **Stop Locating Virtual Drive** from the menu.

9.19 Monitoring Enclosures

When the MegaRAID Storage Manager software is running, you can see the status of all enclosures connected to the server by selecting the **Physical** tab in the left panel. If an enclosure is operating normally, the icon looks like this: . If an enclosure is not functioning normally—for example, if a fan has failed—an orange, yellow, or red circle appears to the right of the icon.

Information about the enclosure appears in the right panel when you select the **Properties** tab on the main menu screen. A graphical display of enclosure information appears when you select the **Graphical View** tab.


The display in the center of the screen shows how many slots of the enclosure are actually populated by the drives and the lights on the drives show the drive status. The information on the right shows you the status of the temperature sensors, fans, and power supplies in the enclosure.

To view the enclosure properties, in the physical view click on the **Enclosure** node. The **Enclosure Properties** are displayed, as shown in the following figure.

Vendor ID	DELL	FRU Number	42R5138
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	PU1100	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16VTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

Figure 257: Enclosure Properties

9.19.1 Monitoring Battery Backup Units

When the MegaRAID Storage Manager software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this . If a BBU fails, a red dot appears next to the icon.

To show the properties for a BBU, perform the following steps:

1. On the main menu screen, click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.

The BBU properties appear in the right panel. The BBU properties include the following:

- The number of times the BBU has been recharged (cycle count).
- The full capacity of the BBU, plus the percentage of its current state of charge, and the estimated time until it will be depleted.
- The current BBU temperature, voltage, current, and remaining capacity.
- If the battery is charging, the estimated time until it is fully charged.
- The battery state, which says if it is in operational state.
- If battery replacement is required.
- The BBU retention time, which gives the total number of hours the battery can support the current capacity reserve.

The BBU Properties are displayed, as shown in the following two figures.

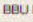
Properties	
BBU Battery Type	iBBU
Battery State	Operational
Battery Replacement	Required / Not required
Temperature	29.0 C (84.2 F) - Normal
Voltage	4065 mV
Current	0 mA
Full Capacity	<value> mAh
Remaining Capacity	<value> mAh
BBU Retention Time	 48+ Hours
Estimated Time to Recharge	<value> Mins
FRU	None
Cycle Count	32
Automatic Learn Cycle	Enabled
Auto Learn Period	30 days
Next Learn Cycle	Aug 10 2010 20:52:13
Relative State of Charge	66%
Absolute State of Charge	35%
Run Time to Empty	Battery is not being discharged
Average Time to Empty	Battery is not being discharged
Average Time to Full	Battery is not being discharged
Maximum Error Margin	25%

Figure 258: Battery Backup Unit Properties for iBBU Battery

Properties			
BBU Battery Type	TMM-C (Not activated) 2	Estimated Time to Recharge	<value> Mins
Battery State	Operational	FRU	None
Battery Replacement	Required / Not required	Memory Module FRU	<value>
Temperature	29.0 C (84.2 F) - Normal	Automatic Learn Cycle	Enabled
Voltage	4055 mV	Auto Learn Period	30 days
Current	0 mA	Next Learn Cycle	Aug 10 2010 20:52:13
Full Capacity	<value> Joules		
Remaining Capacity	<value> Joules		
BBU Retention Time	48+ Hours		

Figure 259: Battery Backup Unit Properties for TMM-C Battery

9.20 Battery Learn Cycle

Learn cycle is a battery calibration operation that is performed by the controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically. To choose automatic battery learn cycles, enable automatic learn cycles. To choose manual battery learn cycles, disable automatic learn cycles.

If you enable automatic learn cycles, you can delay the start of the learn cycles for up to 168 hours (7 days). If you disable automatic learn cycles, you can start the learn cycles manually, and you can choose to receive a reminder to start a manual learn cycle.

9.20.1 Setting Learn Cycle Properties

To set the learn cycle properties, perform the following steps:

1. Click the **Physical** tab to open the Physical view.
2. Select the **BBU** icon in the left panel.
3. Select **Go To->BBU->Set Learn Cycle Properties**. The Set Learn Cycle Properties window appears, as shown in the following figure.

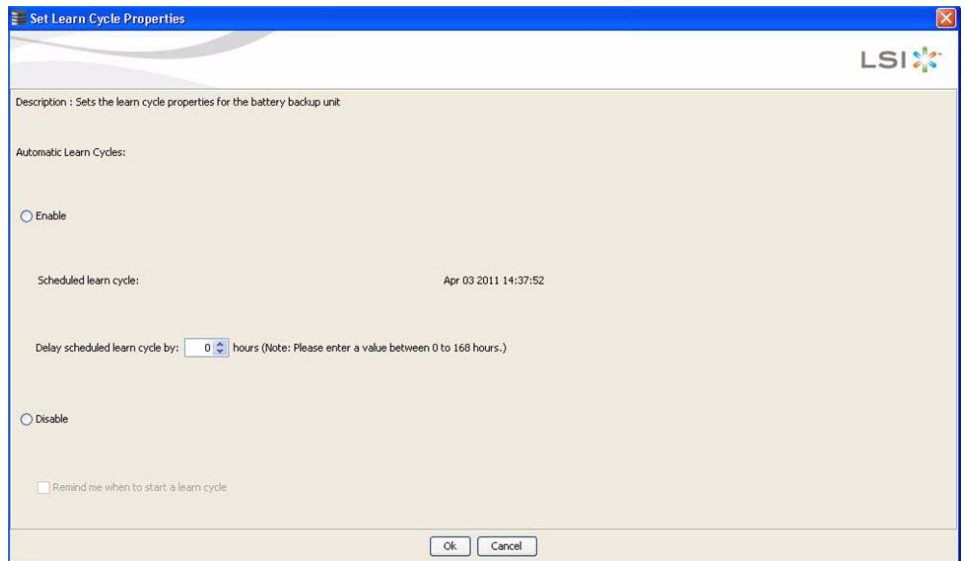


Figure 260: Set Learn Cycle Properties

4. Select one of the two automatic learn cycles:
 - Select the **Enable** radio button to enable an automatic learn cycle.
 - Select the **Disable** radio button to disable an automatic learn cycle.
5. You can delay the start of the next learn cycle up to 7 days (168 hours) by specifying the hours in the **Delay scheduled learn cycle by** field.
6. Select the **Remind me when to start a learn cycle** check box to receive a reminder to start a manual learn cycle.

NOTE: After selecting Disable, if you select Enable, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. The **Next Learn cycle** field (in [Figure 258](#) or [Figure 259](#)) will not be updated until the battery relearn is completed. Once the relearning cycle is completed, the value in the **Next Learn cycle** field will display the new date and the time of the next battery learning cycle.

9.20.2 Starting a Learn Cycle Manually

To start the learn cycle properties manually, perform the following steps:

1. Click the **Physical** tab to open the Physical view.
2. Select the **BBU** icon in the left panel.
3. Select **Go To->BBU->Start Learn Cycle**.

Another way to start the learn cycle is to right-click the **BBU** icon and select **Start Learn Cycle** from the pop-up menu.

9.21 Monitoring Rebuilds and Other Processes

The MegaRAID Storage Manager software allows you to monitor the progress of rebuilds and other lengthy processes in the Group Show Progress window. To open this window, follow these steps:

Follow these steps to monitor the progress of these operations.

Open the show progress window, by selecting **Manage->Show Progress** on the menu bar.

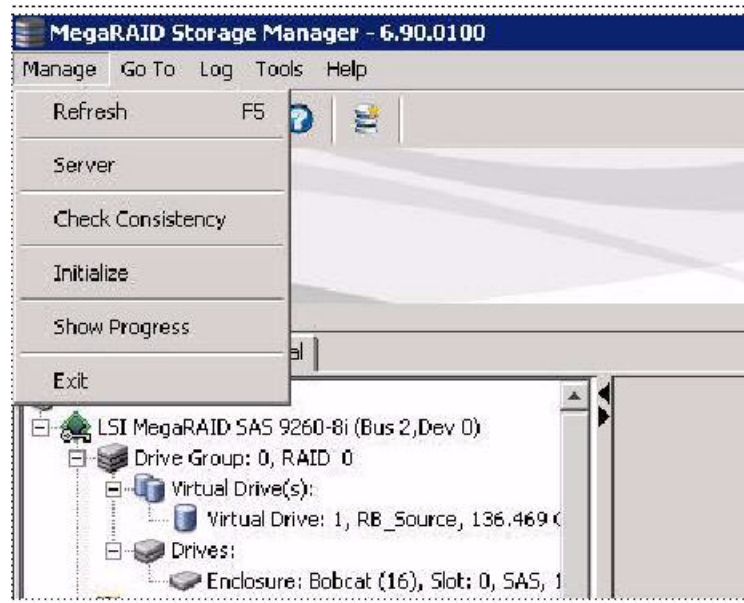


Figure 261: Group Show Progress Menu

The Group Show Progress window appears, as shown in the following figure.

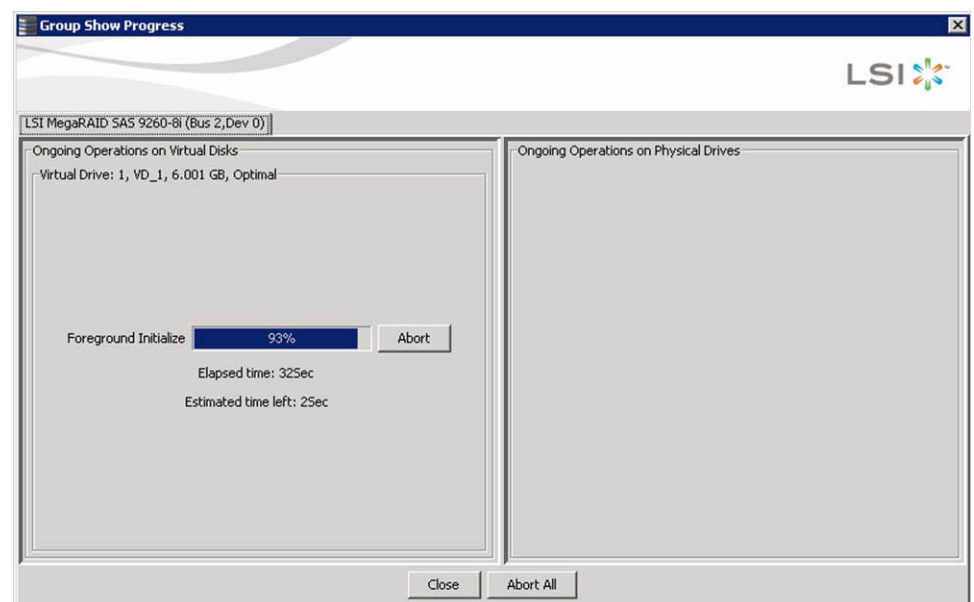


Figure 262: Group Show Progress Window

The Group Show Progress window displays a percent-complete indicator for drive rebuilds. Rebuilds may take a long time to complete and cannot be aborted. An up-arrow appears above the drive icon while it is being rebuilt.

Operations on virtual drives appear in the left panel of the window, and operations on drives appear in the right panel. The type of operations that appear in this window are as follows:

- Initialization of a virtual drive (see [Initializing a Virtual Drive](#))
- Rebuild (see [Rebuilding a Drive](#))
- Consistency check (see [Running a Consistency Check](#))
- Non FDE Physical Drive Erase
- Virtual Drive Erase

A Modify Drive Group process cannot be aborted. To abort any other ongoing process, click the **Abort** button next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

Chapter 10

Maintaining and Managing Storage Configurations

This chapter explains how to use the MegaRAID Storage Manager software to maintain and manage storage configurations. Log on to the server in Full Access mode to perform the maintenance and management tasks.

10.1 Initializing a Virtual Drive

When you create a new virtual drive with the Configuration Wizard, you can select the **Quick Init** or **Full Init** option to initialize the disk immediately. However, you can select **No Init** if you want to initialize the virtual drive later.

To initialize a virtual drive after completing the configuration process, follow these steps:

1. Select the **Logical** tab in the left panel of the MegaRAID Storage Manager window, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To >>Virtual Drive>>Start Initialization**.

The initialize dialog appears.

3. Select the virtual drives to initialize.

CAUTION: Initialization erases all data on the virtual drive. Make sure to back up any data you want to keep before you initialize a virtual drive. Make sure the operating system is not installed on the virtual drive you are initializing.

4. Select the **Fast Initialization** check box if you want to use this option.

If you leave the box unselected, the MegaRAID Storage Manager software will run a Full Initialization on the virtual drive. (For more information, see [Section 8.1.1, Selecting Virtual Drive Settings](#).)

5. Click **Start** to begin the initialization.

You can monitor the progress of the initialization. See [Section 9.21, Monitoring Rebuilds and Other Processes](#) for more information.

10.1.1 Running a Group Initialization

Initialization prepares the storage medium for use. You can run an initialization on multiple drives at one time. Follow these steps to run a group consistency check.

1. Select **Manage>>Initialize**.

The Group Initialization dialog appears, as shown in [Figure 267](#).

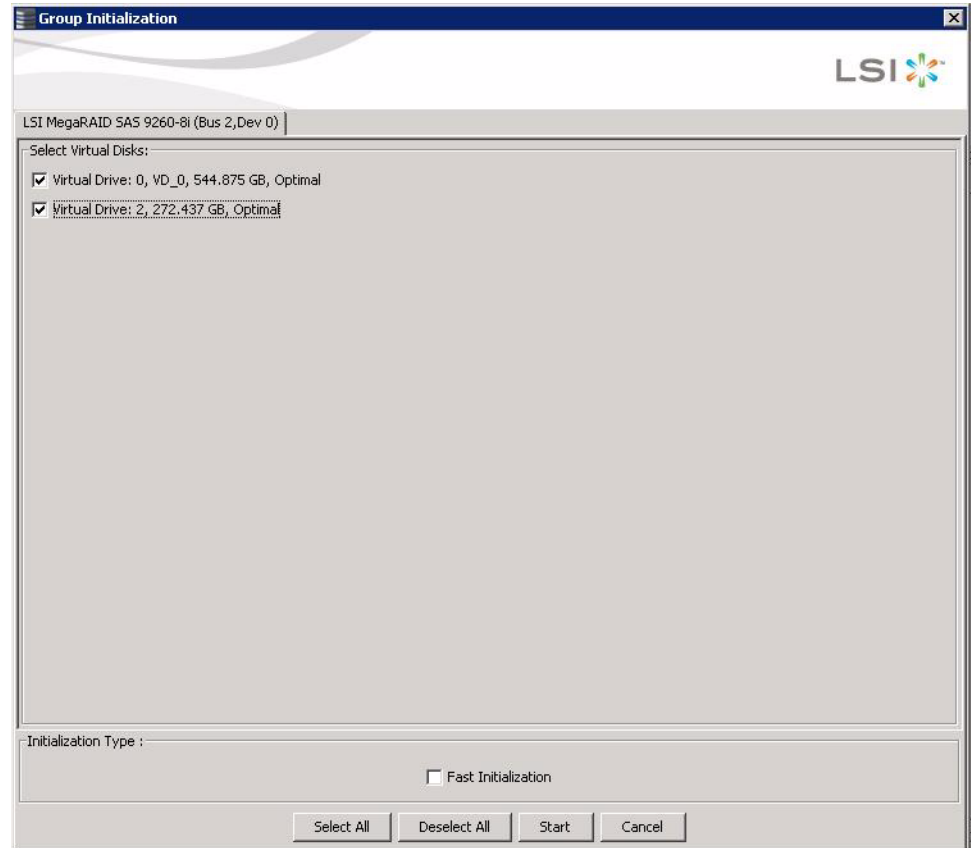


Figure 263: Group Initialization Dialog Box

2. Either check the virtual drives on which to run the initialization, or click **Select All** to select all of the virtual drives.

3. Click **Start**.

You can monitor the progress of the group initialization. See [Section 9.21, Monitoring Rebuilds and Other Processes](#) for more information.

10.2 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives (RAID 1, 5, 6, 10, 50, or 60 configurations; RAID 0 does not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

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. You must run a consistency check if you suspect that the data on the virtual drive might be corrupted.

NOTE: Make sure to back up the data before running a consistency check if you think the data might be corrupted.

To run a consistency check, first set the consistency check properties, and then schedule the consistency check. This section explains how to set the properties, schedule the check, and run the consistency check.

10.2.1 Setting the Consistency Check Settings

Follow these steps to set the properties for a consistency check:

1. Click the **Physical** tab or the **Logical** tab and select a controller.
2. Click **Go To >> Controller >> Set Consistency Check Properties**, as shown in the following figure.

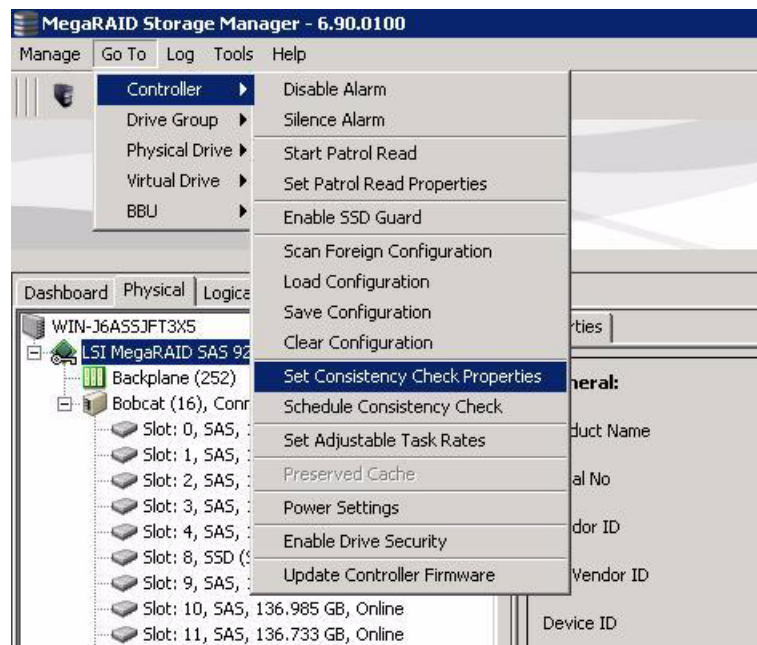


Figure 264: Set Consistency Check Properties Option

The Set Consistency Check Properties dialog appears, as shown in the following figure.

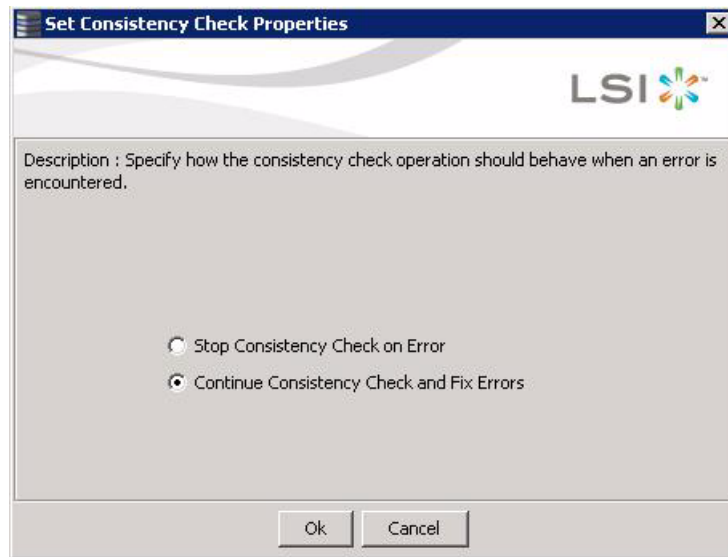


Figure 265: Set Consistency Check Properties Dialog Box

3. Choose one of the two options:
 - **Stop Consistency Check on Error:** The RAID controller stops the consistency check operation if the utility finds an error.
 - **Continue Consistency Check and Fix Errors:** The RAID controller continues the consistency check if the utility finds an error, and then fixes the errors.
4. Click **Ok**.

10.2.2 Scheduling a Consistency Check

Follow these steps to set the properties for a consistency check:

1. Click the **Physical** tab or the **Logical** tab and select the controller.
2. Select **Go To >> Controller >> Schedule Consistency Check**.

The Schedule Consistency Check dialog appears, as shown in the following figure.

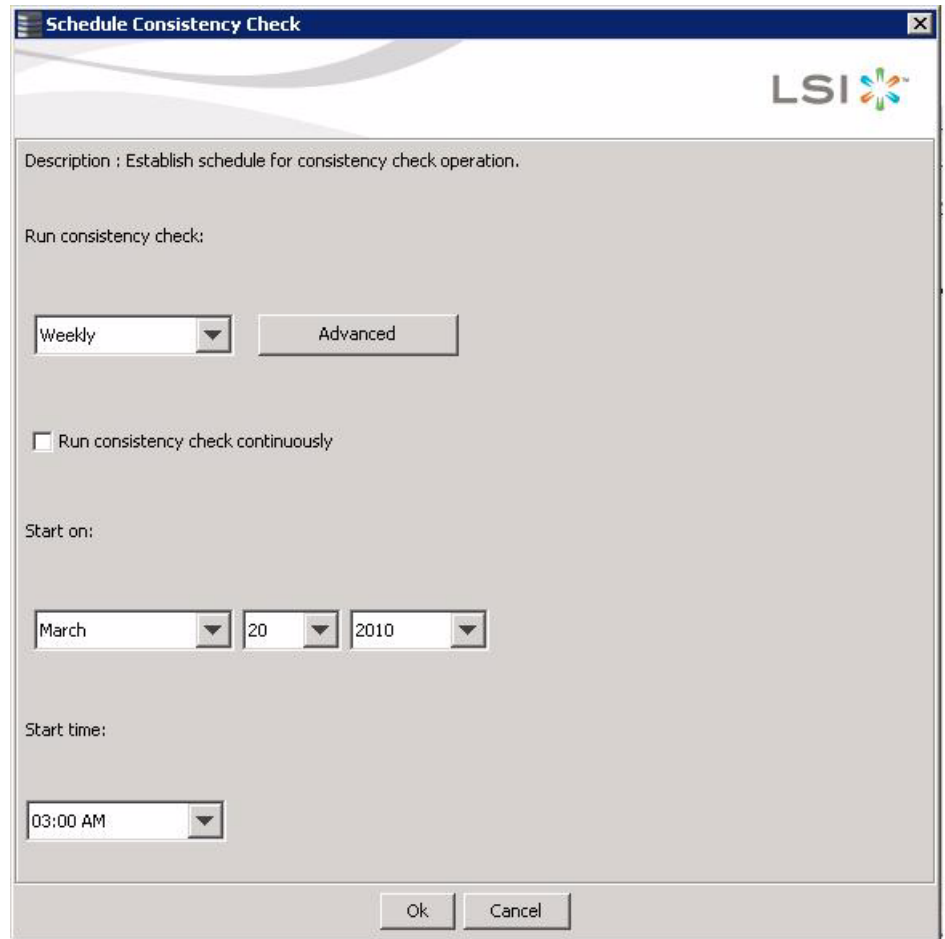


Figure 266: Schedule Consistency Check Dialog

3. Perform the following steps to schedule the consistency check:
 - a. Select how often to run the consistency check from the drop-down list.
You can click **Advanced** for more detailed date options.
 - b. (Optional) Select the **Run consistency check continuously** check box.
 - c. Select the month, day, and year on which to start the consistency check.
 - d. Select the time of day to start the consistency check.
4. Click **Ok**.

You can monitor the progress of the consistency check. See [Section 9.21, Monitoring Rebuilds and Other Processes](#) for more information.

10.2.3 Running a Group Consistency Check

You can run a consistency check on multiple drives at one time. Follow these steps to run a group consistency check.

1. Select **Manage >> Check Consistency**.

The Group Consistency Check dialog appears, as shown in the following figure.

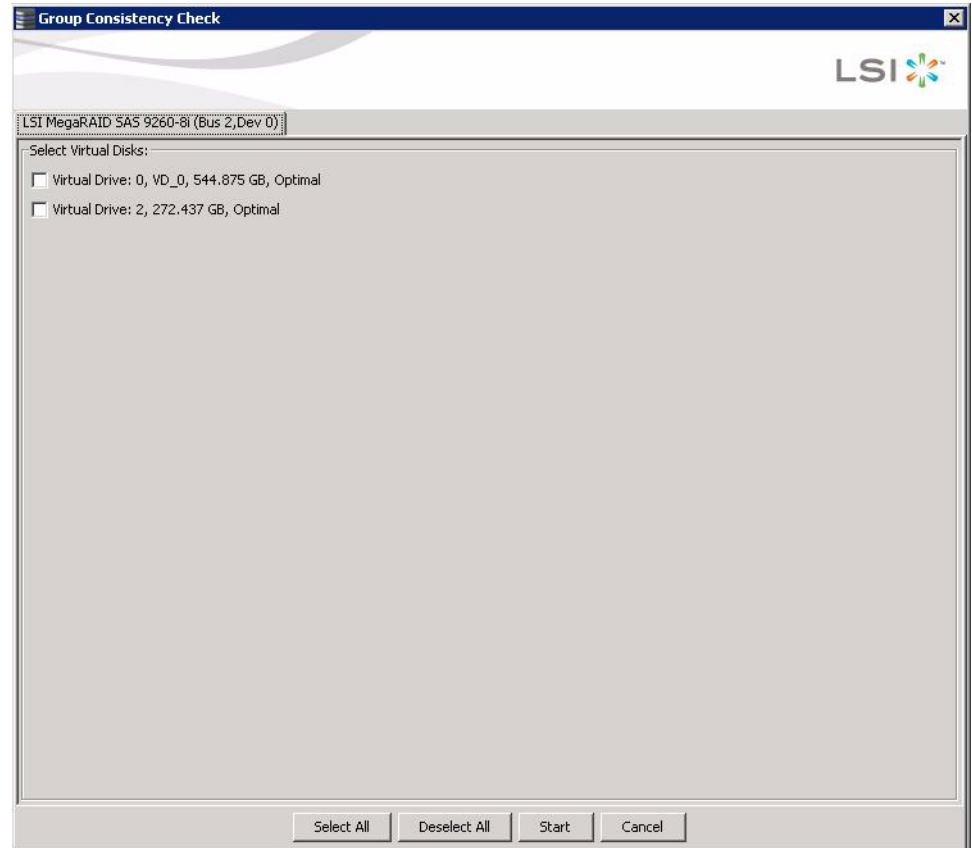


Figure 267: Group Consistency Check Dialog Box

2. Either check the virtual drives on which to run the consistency check, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group consistency check. See [Section 9.21, Monitoring Rebuilds and Other Processes](#) for more information.

10.3 Scanning for New Drives

You can use the Scan for Foreign Configuration option to find drives with foreign configurations. A foreign configuration is a RAID configuration that already exists on a replacement set of physical disks 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. Drives that are foreign are listed on the physical drives list with a special symbol in the MegaRAID Storage Manager software.

The utility allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new configuration using these drives. You can preview the foreign configuration before you decide whether to import it.

The MegaRAID Storage Manager software usually detects newly installed drives and displays icons for them in the MegaRAID Storage Manager window. If for some reason the MegaRAID Storage Manager software does not detect a new drive (or drives), you can use the Scan for Foreign Configuration command to find it.

Follow these steps to scan for a foreign configuration:

1. Select a controller icon in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To >> Controller >> Scan for Foreign Configuration**.

If the MegaRAID Storage Manager software detects any new drives, it displays a list of them on the window. If not, it notifies you that no foreign configuration is found.



3. Follow the instructions on the window to complete the drive detection.

10.4 Rebuilding a Drive

If a drive in a redundant virtual drive (RAID 1, 5, 6, 10, 50, or 60) fails, the MegaRAID Storage Manager software automatically rebuilds the data on a hot spare drive to prevent data loss. *The rebuild is a fully automatic process*, so it is not necessary to issue a Rebuild command. You can monitor the progress of drive rebuilds in the Group Show Progress window. To open this window, select **Group Operations->Show Progress**.

If a single drive in a RAID 1, RAID 5, RAID 10, or RAID 50 virtual drive fails, the system is protected from data loss. A RAID 6 virtual drive can survive two failed drives. A RAID 60 virtual drive can survive two failed drives in each span in the drive group. Data loss is prevented by using parity data in RAID 5, RAID 6, RAID 50, and RAID 60, and data redundancy in RAID 1 and RAID 10.

The failed drive must be replaced, and the data on the drive must be rebuilt on a new drive to restore the system to fault tolerance. You can choose to rebuild the data on the failed drive if the drive is still operational. If dedicated hot spares or global hot spare disks are available, the failed drive is rebuilt automatically without any user intervention.

A red circle to the right of the drive icon   indicates that a drive has failed. A yellow circle appears to the right of the icon of the virtual drive that uses this drive which indicates that the virtual drive is in a degraded state; the data is still safe, but data could be lost if another drive fails.

Follow these steps to rebuild a drive:

1. Right-click the icon of the failed drive, and select **Rebuild**.
2. Click **Yes** when the warning message appears. If the drive is still good, a rebuild will start.

You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Manage >> Show Progress**. If the drive cannot be rebuilt, an error message appears. Continue with the next step.

3. Shut down the system, disconnect the power cord, and open the computer case.
4. Replace the failed drive with a new drive of equal capacity.
5. Close the computer case, reconnect the power cord, and restart the computer.
6. Restart the MegaRAID Storage Manager software.

When the new drive spins up, the drive icon changes back to normal status, and the rebuild process begins automatically. You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Manage >> Show Progress**.

If you want to force a drive into Fail status to trigger a rebuild, right-click the drive icon, and select **Make Drive Offline**. A red circle appears next to the drive icon. Right-click the icon, and select **Rebuild** from the pop-up menu. A drive rebuild cannot be aborted.

NOTE: A drive rebuild is also started if you select **Make Drive Online** from the pop-up menu.

10.4.1 New Drives Attached to a MegaRAID Controller

When you insert a new drive on a MegaRAID system and if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD for MegaRAID entry-level controllers, such as the SAS 9240-4i/8i. If the drive does contain valid DDF metadata, its drive state is Unconfigured Good.

A new drive in JBOD drive state is exposed to the host operating system as a stand-alone drive. Drives in JBOD drive state are not part of the RAID configuration because they do not have valid DDF records. The operating system can install and run anything on JBOD drives.

Automatic rebuilds always occur when the drive slot status changes, for example, when you insert a drive or remove a drive, so that a hot spare can be used. However, a new drive in JBOD drive state (without a valid DDF record), does not perform an automatic rebuild.

To start an automatic rebuild on the new JBOD drive, you have to change the drive state from JBOD to Unconfigured Good. (Rebuilds start on Unconfigured Good drives only.) After you set the drive state to Unconfigured Good, the drive state information always remains on the drive, and you can use the drive for configuration.

10.5 Making a Drive Offline or Missing

If a drive is currently part of a redundant configuration and you want to use it in another configuration, you can use the MegaRAID Storage Manager commands to remove the drive from the first configuration and change the drive state to Unconfigured Good.

CAUTION: After you perform this procedure, *all data on that drive is lost.*

To remove the drive from the configuration without harming the data on the virtual drive, follow these steps:

1. In the MegaRAID Storage Manager window, select **Go To >> Physical Drive >> Make Drive Offline.**

The drive status changes to Offline.

2. Select **Go To >> Physical Drive >> Mark Drive as Missing.**

The drive status changes to Unconfigured Good.

CAUTION: After you perform this step, the data on this drive is no longer valid.

3. If necessary, create a hot spare drive for the virtual drive from which you have removed the drive. (See [Section 8.3, Adding Hot Spare Drives.](#))

When a hot spare is available, the data on the virtual drive will be rebuilt. You can now use the removed drive for another configuration.

CAUTION: If the MegaRAID Storage Manager software detects that a drive in a virtual drive has failed, it makes the drive offline. If this situation occurs, you must remove the drive and replace it. You cannot make the drive usable for another configuration by using the **Mark physical disk as missing** command and the **Rescan** commands.

10.6 Removing a Drive

You may sometimes need to remove a non-failed drive that is connected to the controller. For example, you may need to replace the drive with a larger drive. Follow these steps to remove a drive safely:

1. Click the icon of the drive in the left panel, and click the **Operations** tab in the right panel.
2. Select **Prepare for Removal**, and click **Go**.
3. Wait until the drive spins down and remove it.

If you change your mind, select **Undo Prepare for Removal** and click **Go**.

10.7 Upgrading the Firmware

The MegaRAID Storage Manager software enables you to easily upgrade the controller firmware.

To avoid data loss because of dirty cache on the controller, the utility forces the virtual disks into Write Through mode after a firmware upgrade. It is in this mode until the server reboots. In Write Through mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all of the data in a transaction. This way, in case of a power outage, the controller does not discard the dirty cache.

Follow these steps to upgrade the firmware:

1. In the left panel of the MegaRAID Storage Manager window, click on the icon of the controller you need to upgrade.
2. In the MegaRAID Storage Manager window, select **Go To >> Controller >> Update Controller Firmware**.
3. Click **Browse** to locate the .rom update file, as shown in the following figure.

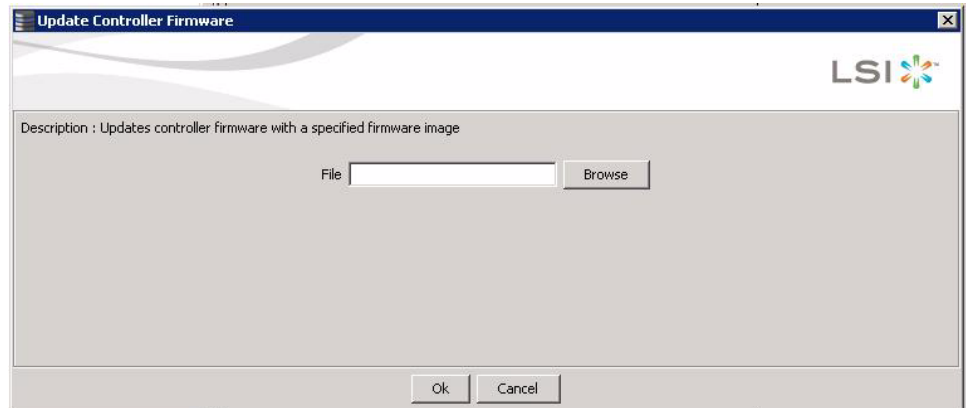


Figure 268: Update Controller Firmware Dialog

4. After you locate the file, click **Ok**.

The MegaRAID Storage Manager software displays the version of the existing firmware and the version of the new firmware file.

5. When you are prompted to indicate whether you want to upgrade the firmware, click **Yes**.

The controller is updated with the new firmware code contained in the .rom file.

6. Reboot the system after the new firmware is flashed.

The new firmware does not take effect until reboot.

Chapter 11

Using MegaRAID Advanced Software

This chapter describes the MegaRAID advanced software offered by the MegaRAID Storage Manager software for certain MegaRAID SAS 6Gb/s RAID controllers and explains how to use these features.

11.1 MegaRAID Advanced Software

The MegaRAID advanced software are features that the MegaRAID Storage Manager software and WebBIOS support on certain MegaRAID SAS 6Gb/s RAID controllers. The following MegaRAID SAS 6Gb/s RAID controllers support advanced software features that offer improved performance, data protection, and availability:

- MegaRAID SAS 9260-4i
- MegaRAID SAS 9260-8i
- MegaRAID SAS 9280-4i4e

NOTE: Record your controller serial number in a safe location in case you need to contact LSI Technical Support.

CAUTION: Back up your data before you make a change in the system configuration. Failure to do so could result in data loss.

11.2 Recovery Advanced Software

The MegaRAID advanced software include the following features.

- MegaRAID FastPath
- MegaRAID Recovery
- MegaRAID CacheCade 2.0 SSD Read Caching software
- MegaRAID CacheCade Pro 2.0 SSD Read/Write Caching software
- MegaRAID RAID 6
- MegaRAID RAID 5

11.2.1 MegaRAID Software Licensing

The MegaRAID Software licensing authorizes you to enable the MegaRAID advanced software features present in the MegaRAID Storage Manager application. You have to obtain the activation key to enable, and use the advanced software features present in the controller. You can also implement the rehosting process by configuring the key vault, if you want to transfer the advanced features from one controller to another.

11.2.2 Managing MegaRAID Advanced Software

The MegaRAID Advanced Software wizard allows you to use the advanced software features. Perform the following steps to enable the *activation key* to use the advanced controller features:

1. Select the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.
2. Choose either of the following options:
 - Select **Go To>>Controller >>Manage MegaRAID Advanced Software Options**,
 - Click **Manage MegaRAID Advanced Software Options** from the dashboard under the feature portlet.

The Manage MegaRAID Advanced Software Options wizard appears.

- If none of the advanced software options present in the controller are in a boot mode, [Figure 269](#) appears.
- If even one of the advanced software options present in the controller is in a boot mode, [Figure 270](#) appears. You cannot activate any advanced software option from this window as this is a view-only window.

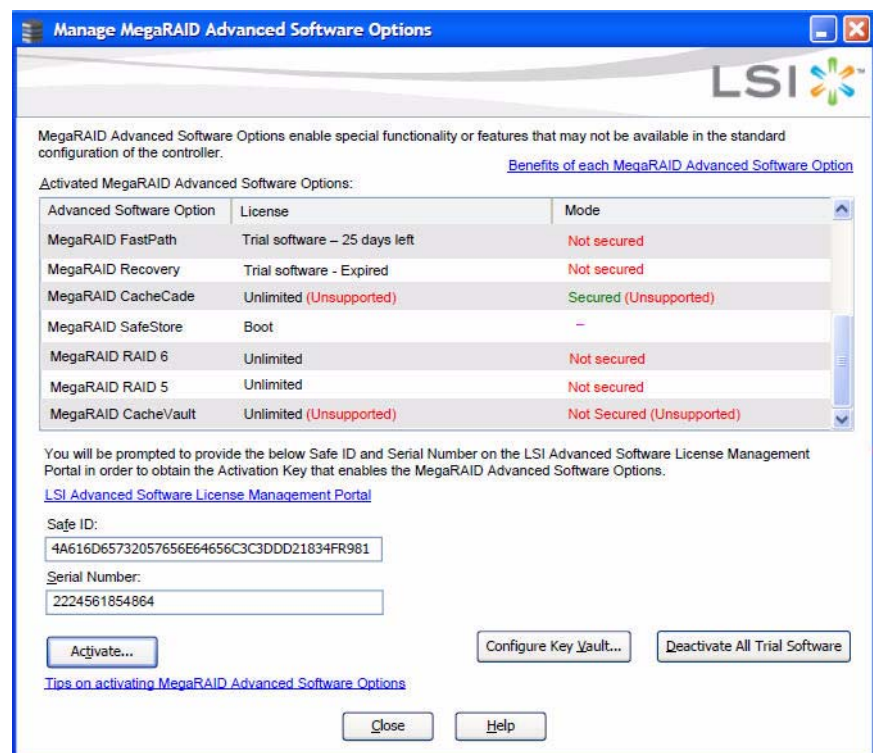


Figure 269: Manage MegaRAID Advanced Software Options Dialog

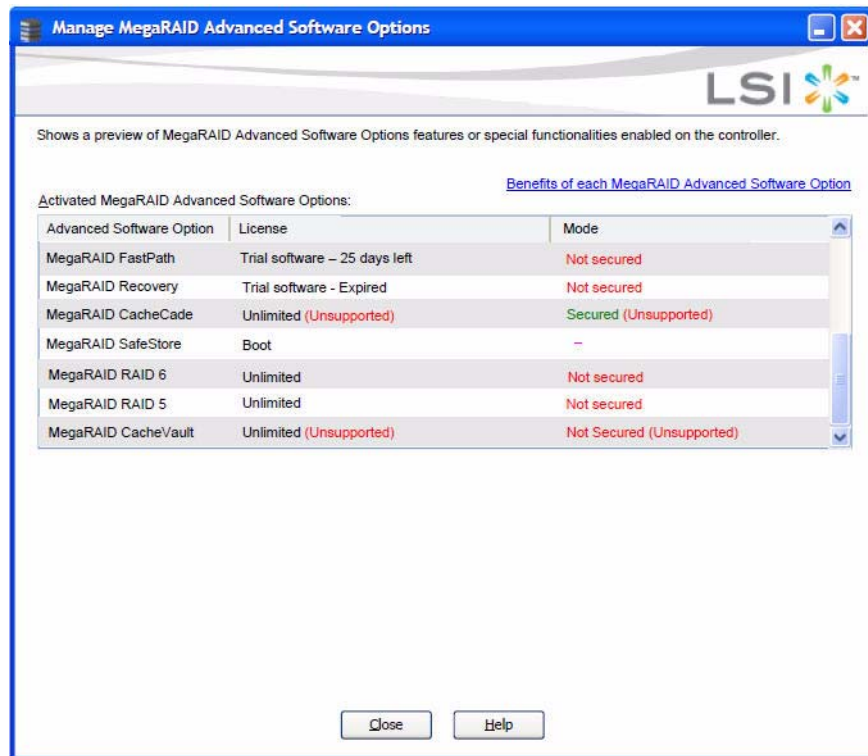


Figure 270: Manage MegaRAID Advanced Software Options Dialog

The **Activated MegaRAID Advanced Software Options** table consists of the **Advanced Software Option**, **License**, and **Mode** columns.

- The **Advanced Software Option** column displays the list of advanced software options present in the controller.

NOTE: In the above figure, instead of the option **MegaRAID CacheCade**, the options **MegaRAID CacheCade 2.0 software** and **MegaRAID CacheCade Pro 2.0 software** will be displayed.

- The **License** column displays the license details for the list of advanced software options present in the **Advanced Software Option** column. The license details validates if the software is under trial period, or if it can be used without any trial period (Unlimited).
- The **Mode** column displays the current status of the advanced software. The current status can be Secured, Not secured, or Factory installed.

NOTE: The **Mode** column appears only if the Key Vault is present.

3. Click the **LSI Advanced Software License Management Portal** link to obtain the license authorization code and activation key.

If you click the **Benefits of each MegaRAID Advanced Software** link, you can access http://www.lsi.com/channel/products/advanced_software. If you click the **Tips on activating MegaRAID Advanced Software** link, you can access www.lsi.com/channel/licensing.

Both the **Safe ID** field and the **Serial Number** field consists of a pre-defined value generated by the controller. Alternatively, you can copy the value and paste it in the text box for the applicable field.

4. Click **Activate**.

The Activate MegaRAID Advanced Software – Choose Method wizard appears, as shown in the following figure.

11.2.3 Activation Key

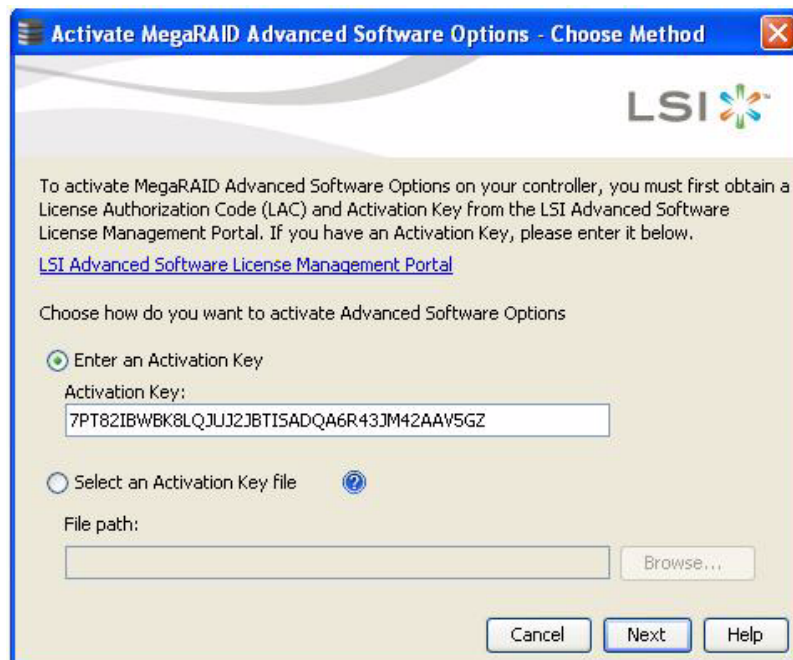


Figure 271: Activate MegaRAID Advanced Software Options– Choose Method Dialog

Perform the following steps to enter the activation key:

1. Click the **LSI Advanced Software License Management Portal** link to obtain a license authorization code (LAC) and activation key.
2. Use any one of the following options to enter the activation key:
 - a. Select the **Enter an Activation Key** radio button, and enter the activation key in the text box provided below the **Activation Key** field.
 - b. Select the **Select an Activation Key file** radio button, and click **Browse** to get the path of the activation key file.

3. Click **Next**.

After you click **Next**, one of the following two scenarios occurs:

- The Activate MegaRAID Advanced Software Options – Summary dialog appears as shown in [Figure 272](#).
- Depending on the relevant scenarios, the application responds by displaying corresponding messages as shown in [Application Scenarios and Messages](#).

11.2.4 Advanced MegaRAID Software Status Summary

After you enter the activation key, and click **Next**, the Activate MegaRAID Advanced Software Option – Summary wizard (as shown in the following figure) displays the list of the advanced softwares along with their *former status* and *new status* in the controller.

- The **Advanced Software Option** column displays the currently available software in the controller.
- The **Former Status** column displays the status of the available advanced software prior to entering the activation key.
- The **New Status** column displays the status of the available advanced software, after entering the activation key.

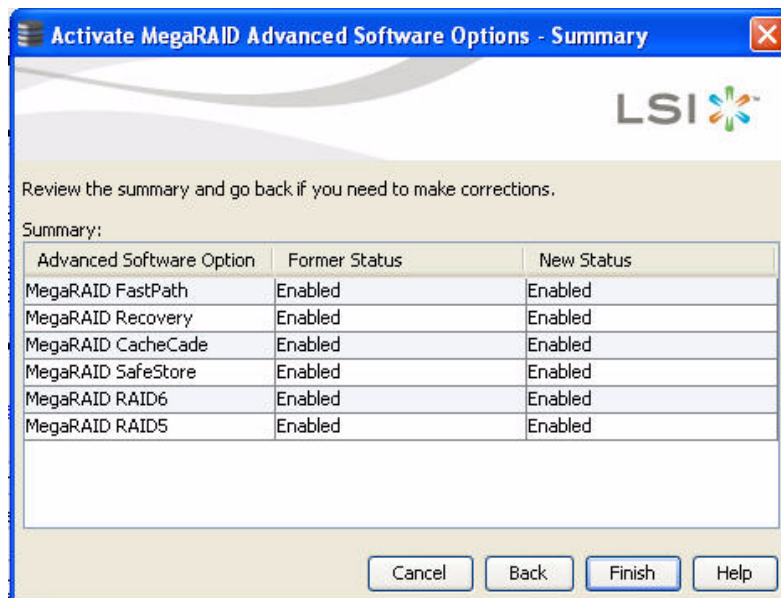


Figure 272: Activate MegaRAID Advanced Software Option– Summary Dialog

4. Click **Finish**.

The status of the advanced software is enabled, and the advanced features are secured in the Key Vault.

5. Click **Cancel** to cancel this action.

NOTE: In [Figure 272](#), instead of the option **MegaRAID CacheCade**, the options **MegaRAID CacheCade 2.0 software** and **MegaRAID CacheCade Pro 2.0 software** will be displayed.

11.2.5 Application Scenarios and Messages

■ Scenario # 1

If you enter an *invalid* activation key, the following message appears.



Figure 273: Invalid Activation Key Message

■ Scenario # 2

If you enter an *incorrect* activation key file, the following message appears.




Figure 274: Incorrect Activation Key File Message

■ Scenario # 3

If you enter an *incorrect* activation key, and if a mismatch exists between the activation key and the controller, the following message appears.




Figure 275: Activation Key Mismatch Message

NOTE: Entering a space in the **Activation Key** field disables the **Next** button in [Figure 271](#).

If you click **Cancel** in the **Activate MegaRAID Advanced Software – Choose Method dialog**, as shown in [Figure 271](#), the following confirmation dialog box appears.

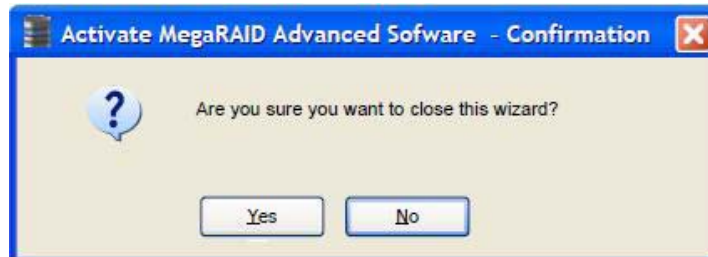


Figure 276: Activate MegaRAID Advanced Software – Confirmation Dialog

11.2.6 Activating an Unlimited Key Over a Trial Key

When you activate an unlimited key over a trial key, a message, The existing trial key will be deactivated and all the advanced software associated with it will be disabled, appears (indicated in pink text in the following figure).

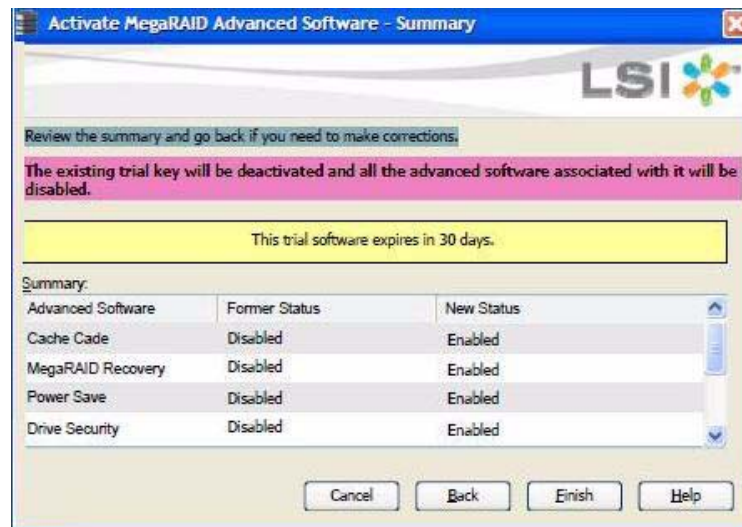


Figure 277: Activating an Unlimited Key over a Trial Key

NOTE: Except for the yellow shading, the other shadings of the text are provided for easy understanding in the relevant dialogs. In the above figure, instead of the option **CacheCade**, the options **MegaRAID CacheCade 2.0 software** and **MegaRAID CacheCade Pro 2.0 software** will be displayed

11.2.6.1 Activating a Trial Software

When you activate a trial software, a message This trial software expires in 30 days appears (indicated in yellow text in the following figure).

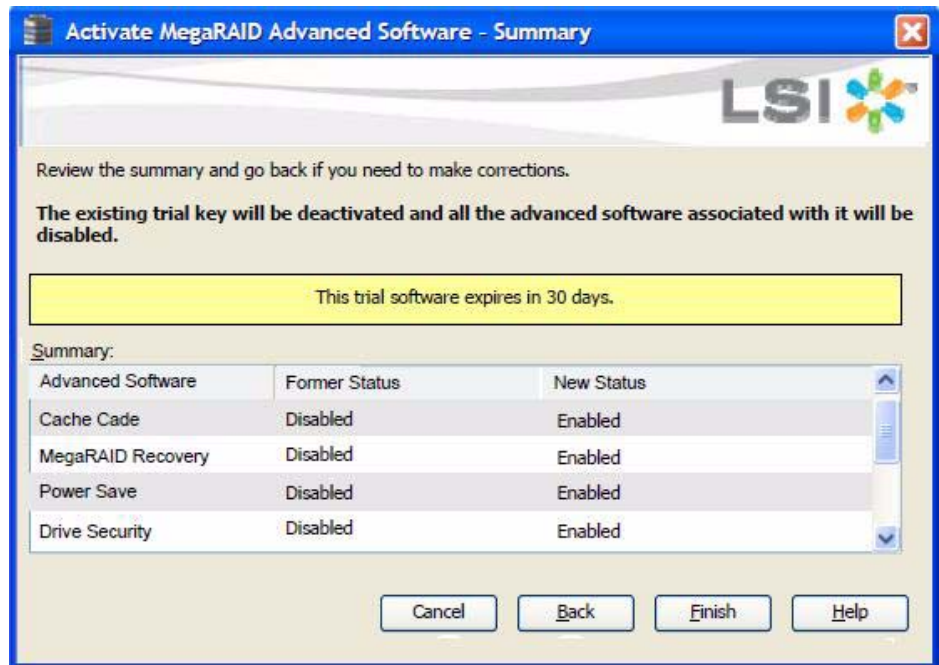


Figure 278: Activating a Trial Software

NOTE: In the above figure, instead of the option **CacheCade**, the options **MegaRAID CacheCade 2.0 software** and **MegaRAID CacheCade Pro 2.0 software** will be displayed.

11.2.6.2 Activating the Unlimited Key

When you activate the unlimited key or a trial key, a message Review the summary and go back if you need to make corrections appears (indicated in green text in the following figure).

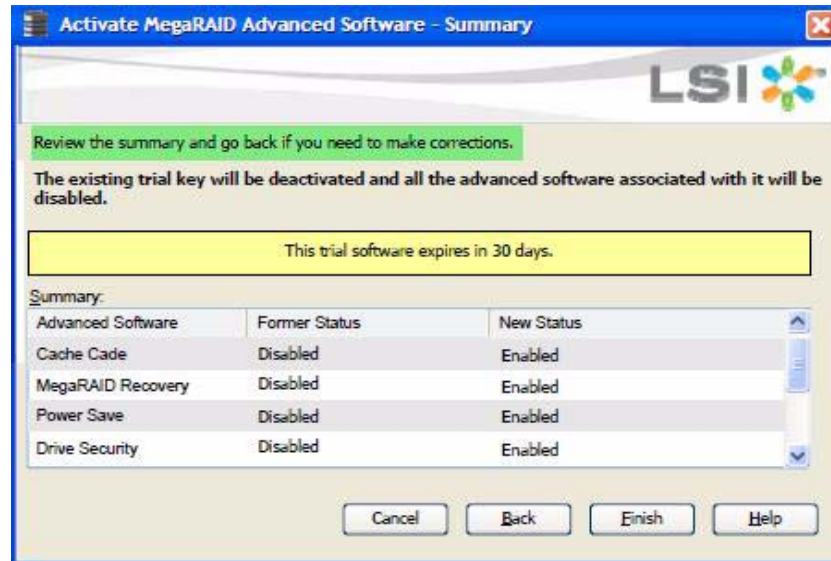


Figure 279: Activating an Unlimited Key

NOTE: In the above figure, instead of the option **CacheCade**, the options **MegaRAID CacheCade 2.0 software** and **MegaRAID CacheCade Pro 2.0 software** will be displayed.

11.2.6.3 Reusing the Activation Key

If you are using an existing activated key, the features are transferred to the key vault, and a message appears, as shown in the following figure.

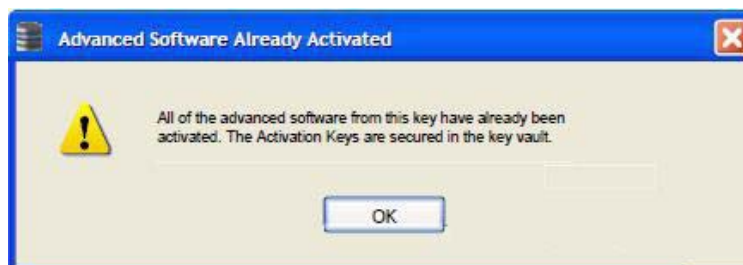


Figure 280: Reusing the Activation Key

11.2.6.4 Securing Advanced MegaRAID Software

When you want to transfer the advanced software from the controller to the Key Vault, use the Securing Advanced MegaRAID Software - Confirmation wizard. This wizard is conditional, and appears only when the Key Vault and the unsecured keys exist.

1. Select any one of the following options to view the Securing Advanced MegaRAID Software - Confirmation wizard.
 - Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.

- Select **Go To>>Controller>>Manage MegaRAID Advanced Software Options** wizard.



Figure 281: Secure MegaRAID Advanced Software - Confirmation Dialog

2. Select the **Confirm** check box, if you want to secure the advanced software.

After you select the check box, the **Yes** button is enabled. This situation implies that the advanced software is secured in the keyvault.

If the advanced software is not secured, the Secure MegaRAID Advanced Software - Confirmation dialog appears, as shown in [Figure 276](#).

11.2.7 Configuring Key Vault (Re-hosting process)

Re-hosting is a process of transferring the advanced software features from one controller to another. To implement the re-hosting process, you must configure the **Configure Key Vault** button in the Manage MegaRAID Software Options wizard, as shown in [Figure 269](#).

Choose any one of the following options to configure the Key Vault.

- Click the Configure Key Vault button in the Manage MegaRAID Advanced Software Options wizard in [Figure 269](#).
- Select **Go To>>Controller >>Manage Premium Feature**.

The Configure Key Vault-Confirm Re-hosting Process wizard appears, as shown in the following figure.

Configure Key Vault - Confirm Re-hosting Process

LSI

To transfer Advanced Software Options from one controller to another controller you need to complete the re-hosting process. Only then you will be able to secure the Advanced Software Options in the key vault.

This wizard helps you to configure the key vault by transferring the Advanced Software Options from one controller to another controller and securing them in the key vault.

Please furnish the below details in the LSI Advanced Software License Management Portal in order to complete the re-hosting process. If you have already completed the process then select the checkbox below and proceed with next.

[LSI Advanced Software License Management Portal](#)

Former Serial Number:

New Serial Number:

Safe ID:

I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal.

Cancel Next Help

Figure 282: Configure Key Vault

3. Select the **I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal** check box.
4. Click **Next**.

The **Configure key Vault- Secure Advanced Software** Options wizard appears, as shown in the following figure.

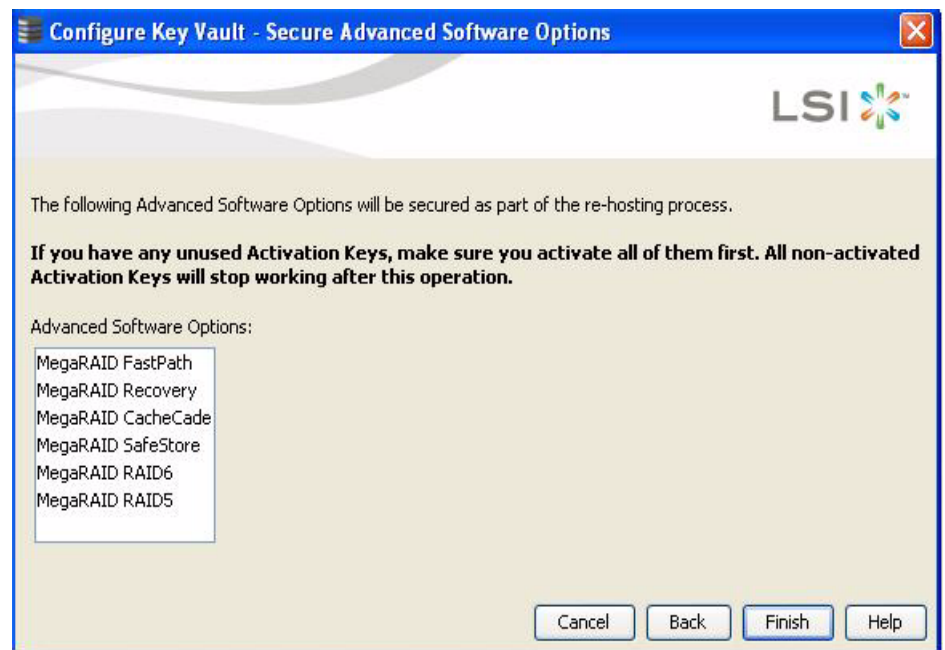


Figure 283: Configure Key Vault - Secure Advanced Software Options Dialog

5. Click **Finish** and the advanced software options are secured in the key vault.

NOTE: The **Next** button in the Configure Key Vault wizard is enabled only if you select the check box.

NOTE: This wizard is conditional and appears only if the re-hosting process is necessary, and when both the key vault and the unsecured keys are present at the same time.

11.2.8 Re-hosting Complete

If you want to transfer the advanced software options from one controller to another, use the re-hosting process. The re-hosting process makes sure that these options are secured in the Key Vault. You have to configure the Key Vault to complete the re-hosting process.

Choose any one of the following options to complete the re-hosting process.

- Click the **Configure Key Vault** button from the Manage MegaRAID Advanced Software Options wizard.
- Select **Go To>>Controller >>Manage MegaRAID Advanced Software Options** wizard.

The **Re-Hosting Process - Complete** wizard appears, as shown in the following figure.

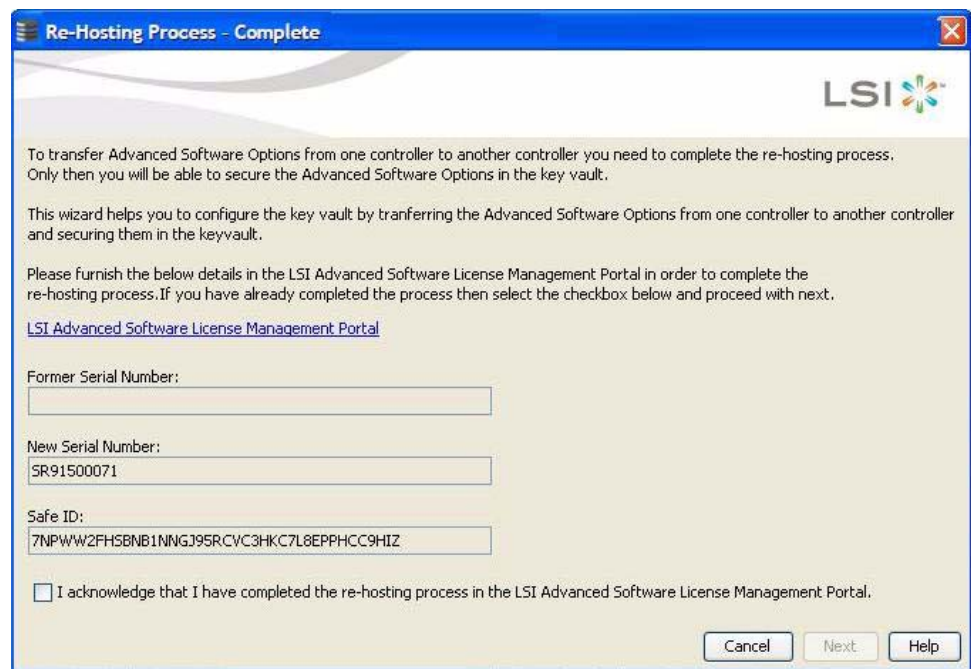


Figure 284: Re-Hosting Process - Complete Dialog

6. Select the **I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal** check box if you want to complete the re-hosting process.

This setting makes sure that the advanced software features are transferred to the controller.

7. Click **Cancel** if you do not want to activate the re-hosting process.

11.2.9 Deactivate Trial Software

When you want to deactivate a trial software, use the **Deactivate All Trial Software** wizard.

Perform the following steps to enable the deactivate trial software button:

1. Click **Deactivate All Trial Software** in the Manage MegaRAID Advanced Software Options wizard as shown in [Figure 269](#).

The Deactivate All Trial Software - Confirmation dialog appears, as shown in the following figure.

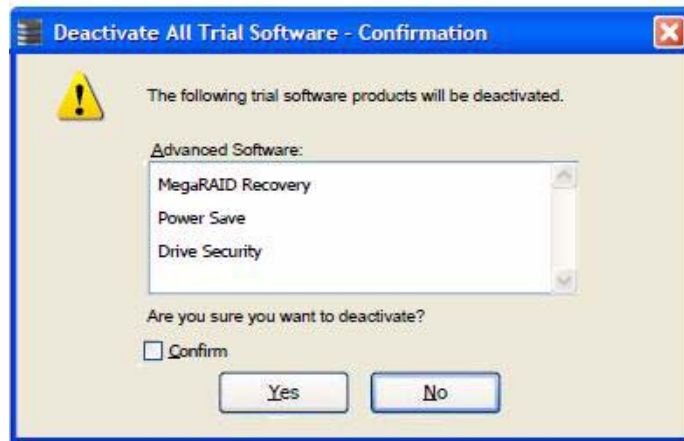


Figure 285: Deactivate All Trial Software – Confirmation Dialog

2. Select the **Confirm** check box, if you want to deactivate the software applications, that are used with a trial key.
3. Click **Yes**.

The trial software is deactivated.

11.2.10 MegaRAID Recovery

MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if the data is deleted accidentally or maliciously, you can restore the data from the view or roll back to a snapshot at a previous point-in-time (PiT). MegaRAID Recovery supports up to 8 snapshots of PiTs for each volume.

Each Recovery PiT volume snapshot is typically a fraction of the original volume size, because it tracks only the changes that are made to a volume after the PiT is created. Disk space for PiTs is reserved in the snapshot repository virtual drive, and the PiT is expanded in small increments as new data is written to the volume. Multiple PiTs of each volume can be retained online, enabling frequent snapshots to be stored in a space-efficient manner.

11.2.11 Recovery Scenarios

There are three primary scenarios in which to use the Recovery feature:

1. Restore the missing or deleted files (restore from view) with the following steps:
 - a. Discover which file is missing or corrupted.
 - b. Review the Snapshot views of the file content (also known as *mounting* the snapshot) from each PiT until you find an earlier version of the missing or corrupted file. A mounted view appears as another drive letter in the Windows Explorer window.
 - c. Drag and drop the earlier version of the file from the Snapshot view back into the online storage volume that was the source of the snapshot.

2. If there is a corrupt volume or operating system, roll back the volume to a previous state with the following steps:
 - a. Restart the system, and press **Ctrl+H** during the power-on self-test (POST).
 - b. In the WebBIOS window, select the corrupted virtual drive, and, on the next dialog that appears, select the **Adv Opers** option.
 - c. Select **Rollback**, and designate the most recent PiT from the drop-down list.
 - d. Click **Go**, and exit WebBIOS.
The system reboots.
 - e. Begin debug and verification procedures on the volume.
You can follow these same steps to roll back to previous PiTs.
3. Reduce the risk of extended downtime during the application updates and upgrades in the IT center with the following steps:
 - a. When the application is offline, take a snapshot of the application volume.
 - b. Install each patch individually, and test for any new defects that might have been introduced.
 - c. Take a snapshot after you test each patch, and determine that it is clean.
 - d. If a defect is introduced, roll back to the previous installation, and bypass the installation of the defective patch.

NOTE: If the volume is still damaged, continue to select from the next most current PiT to the oldest.

11.2.12 Enabling the Recovery Advanced Software

You can enable the Recovery advanced software in the MegaRAID Storage Manager software. When you enable Recovery, you create two virtual drives, one as a snapshot base or a source and the other as a snapshot repository. The base virtual drive contains the data that is stored in the repository virtual drive.

Perform the following these steps to enable MegaRAID recovery:

1. Select the **Logical** tab on the main menu dialog for the Logical view.
2. Select and highlight a virtual drive from the list of virtual drives.

This is the snapshot base virtual drive.

NOTE: A base virtual drive and a repository virtual drive can be associated with the same drives or a common set of drives, or the two virtual drives can be located on two completely separate set of drives. Using a separate set of drives for the base virtual drive and the repository virtual drives provides a performance advantage over using a common set of drives.

3. Select **Go To>>Virtual Drive>>Enable MegaRAID Recovery** on the menu bar, as shown in the following figure.

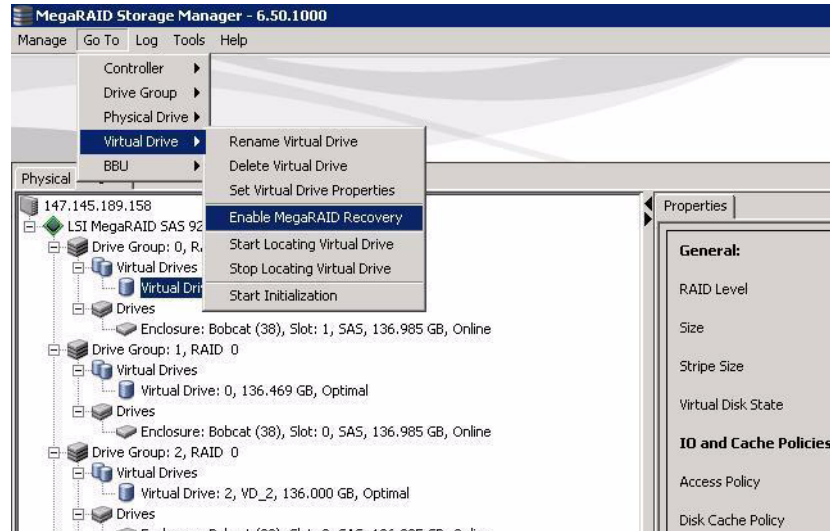


Figure 286: Enable MegaRAID Recovery Menu Option

The Enable MegaRAID Recovery wizard appears, as shown in [Figure 287](#). This wizard allows you to select the virtual drive to use as the snapshot repository.

11.2.13 Snapshot Repository

You can select an existing virtual drive, or create a new virtual drive for the snapshot repository.

Select any one of the options to select or create a new virtual drive.

- Select the virtual drive to use as the snapshot repository in the **Snapshot Repository** field from the **Enable MegaRAID Recovery** wizard as shown in [Figure 287](#).
- Click **Create New** to create a new virtual drive to use as the Snapshot Repository.

When you create a new virtual drive, the newly created virtual drive is appended to the existing rows in the **Snapshot Repository** field.

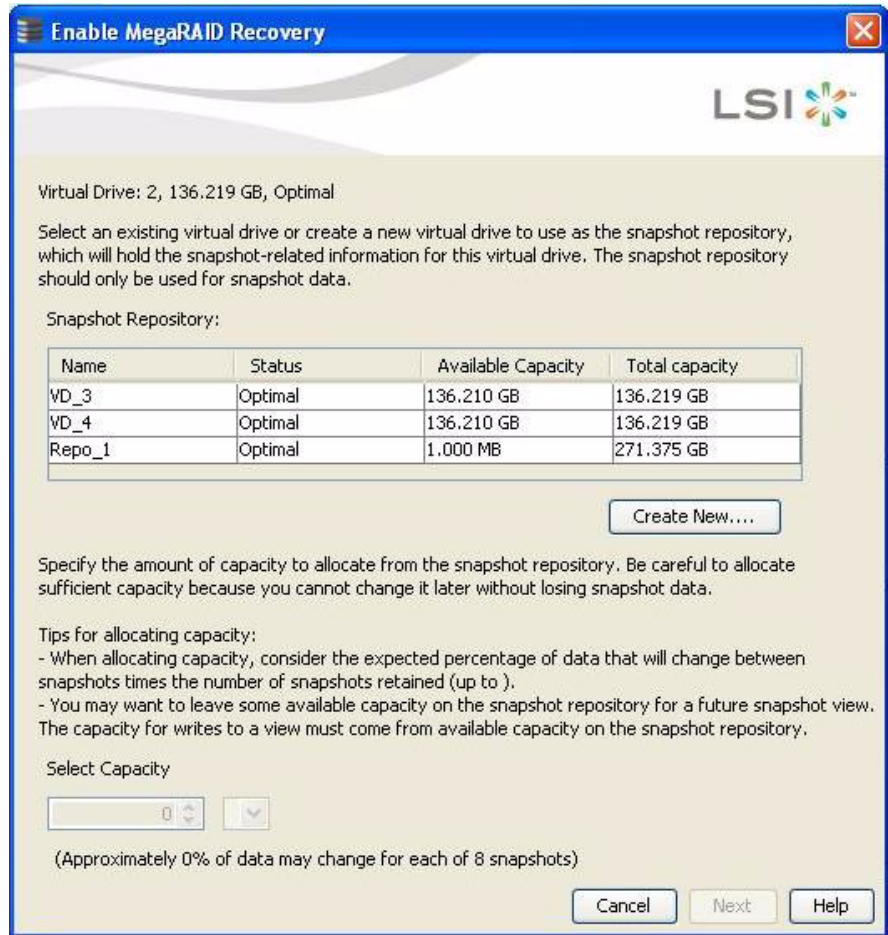


Figure 287: Enable MegaRAID Recovery Wizard

The **Snapshot Repository** table displays the details of the default virtual drives existing in the system, if there are any.

- **Name** – Displays the name of the virtual drive.
- **Status** – Displays the status of the virtual drive.
- **Available Capacity** – Displays the available capacity on the virtual drive.
- **Total Capacity** – Displays the total capacity of the virtual drive.

If the default virtual drives do not exist in the system, the columns in the **Snapshot Repository** table are blank.

The status of the virtual drive can be optimal, degraded, or partially degraded.

4. In the **Select Capacity** field, use the drop-down list to select the appropriate capacity to use for changes to the base virtual drive.

The capacity depends on how write-intensive the application snapshots are. The available capacity is the largest free block of capacity on the snapshot repository virtual drive.

NOTE: Refer to the tips provided for allocating capacity in [Figure 287](#).

NOTE: If you designate all of the capacity for the virtual drive repository, you cannot use the same virtual drive as a repository for other volumes.

5. Click **Next**.

The Enable MegaRAID Recovery - Displaying the Selected Virtual Drive appears as shown in [Figure 288](#). This wizard lets you to select the virtual drive to be used as the snapshot repository.

11.2.14 Selecting the Virtual Drive

You can select the virtual drive to use as the snapshot repository, and also allocate the capacity for the virtual drive from the snapshot repository.

Perform these steps to select the virtual drive in the **Snapshot Repository** field:

1. Select the virtual drive to be used as the snapshot repository to hold the snapshot information. The selected virtual drive is highlighted, as shown in the following figure.

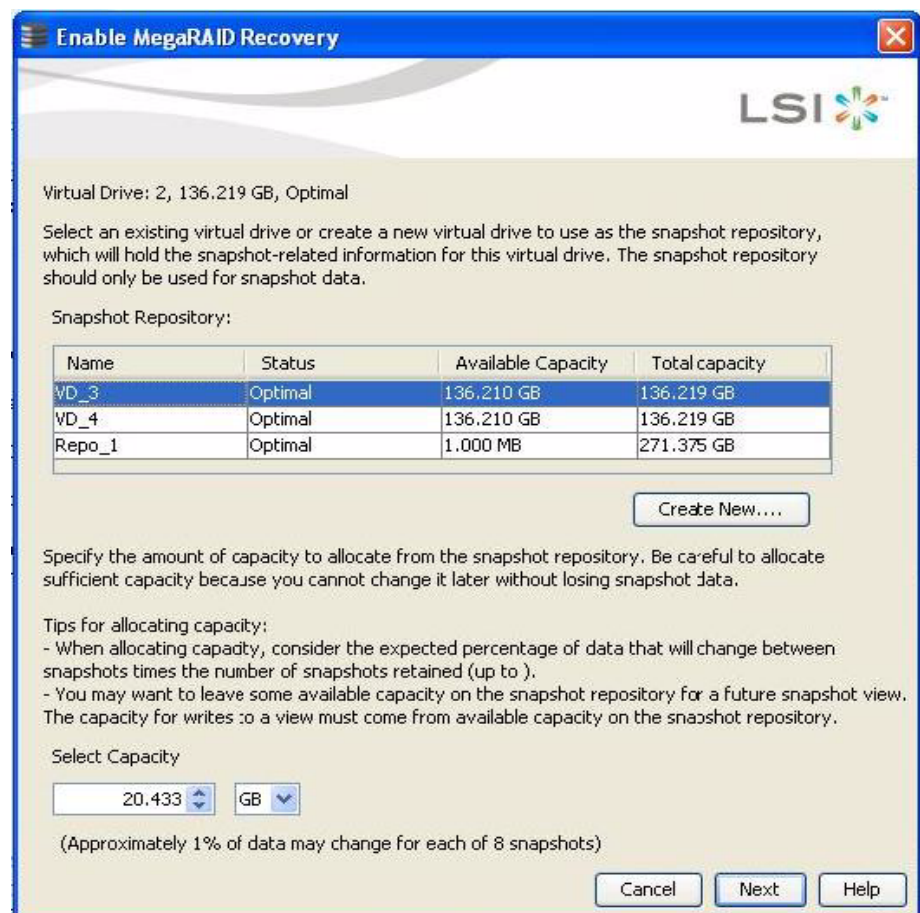


Figure 288: Enable MegaRAID Recovery Wizard – Displaying the Selected Virtual Drive

2. In the **Select Capacity** field, use the drop-down selector to select the appropriate capacity to use for changes to the base virtual drive.
3. Click **Next**.

The Enable MegaRAID Recovery – Create Snapshot Schedule wizard appears, as shown in [Figure 289](#). This wizard lets you to schedule the snapshots.

11.2.15 Scheduling Snapshots

You can select an existing snapshot schedule or create a new snapshot schedule for the virtual drive.

Follow these steps to schedule snapshots.

1. Select any one of the options shown in the Enable MegaRAID Recovery - Create Snapshot Schedule wizard, as shown in the following figure, to schedule snapshots.

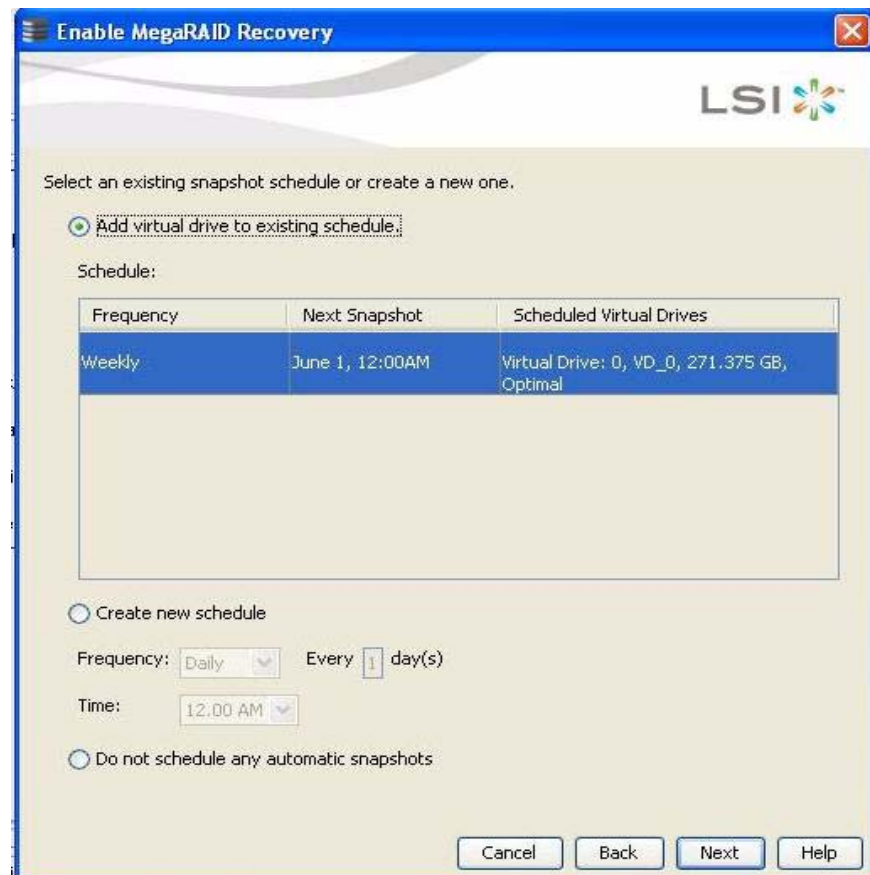


Figure 289: Enable MegaRAID Recovery - Create Snapshot Schedule Dialog

- **Add virtual drive to the existing schedule** — This option allows you to add a virtual drive to a pre-existing schedule already defined in the system.

The **Schedule** table displays the schedule details of the virtual drive.

- The **Frequency** column – Displays the frequency of the daily or weekly snapshot schedule.
- The **Next Snapshot** column – Displays the date and time of the next scheduled snapshot.
- The **Scheduled Virtual Drives** column – Represents the details of the default VDs present in the system.

- **Create new schedule** – This option allows you to create a new schedule to the virtual drive.

- In the **Frequency** field, use the drop-down selector to select the frequency of the snapshot (daily or weekly).
- In the **Time** field, use the drop-down selector to select the time of the scheduled snapshot.

- **Do not schedule any automatic snapshots** – This option prevents you from capturing automatic snapshots from the system.

2. Click **Next**.

The Enable MegaRAID Recovery – Editing Snapshot Properties wizard appears, as shown in [Figure 290](#). You can edit the settings for automatic snapshots.

11.2.16 Editing Snapshots

You can edit the property settings that are already defined for the automatic snapshots using the options as shown in the following figure.

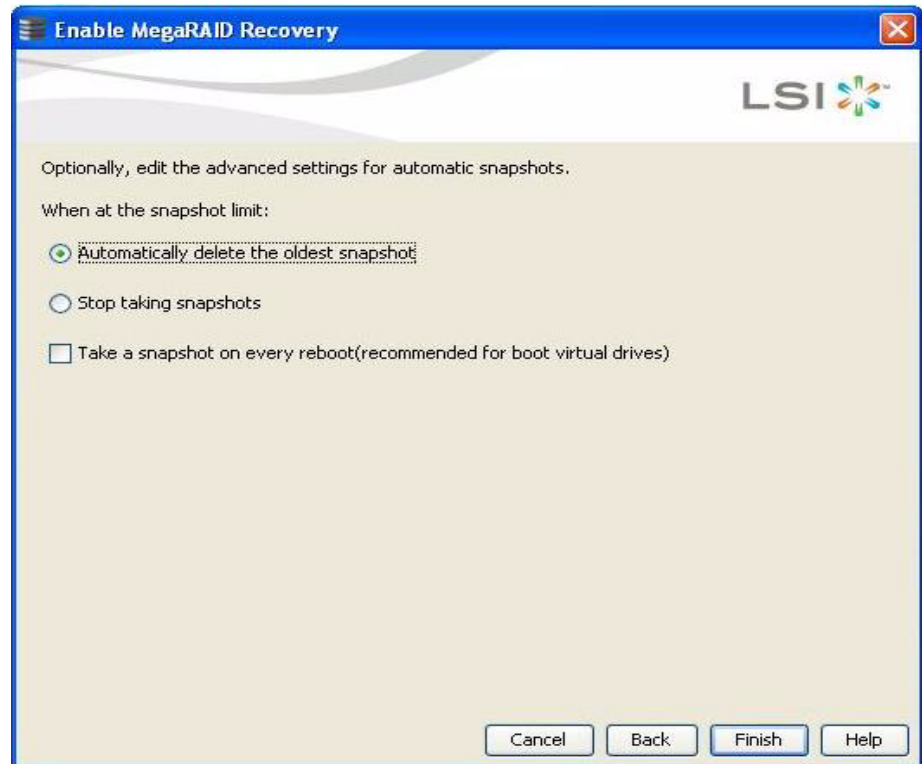


Figure 290: Enable MegaRAID Recovery – Editing Snapshot Properties Dialog

1. Select any one of the options to edit the snapshots.
 - **Automatically delete the oldest snapshot** – This option automatically deletes the oldest snapshot present in the system.
 - **Stop taking snapshots** – This option, prevents the application from taking the snapshots.
 - **Take a snapshot on every reboot(recommended for boot virtual drives)** – To use this option, select the **Take a snapshot on every reboot (recommended for boot virtual drives)** check box.

This option provides you a snapshot taken on boot after each successful shutdown. You can use this snapshot of the boot virtual drive to restore the operating system on the virtual drive if it becomes corrupted.

2. Click **Finish**.

The Confirm Enable Snapshots dialog appears. This dialog prompts you to make sure whether you want to enable snapshots on the virtual drive or not.



Figure 291: Confirm Enable Snapshots

3. If your answer is yes, select the **Confirm** check box.

When you select the **Confirm** check box, the **Yes** button is enabled. The snapshots are enabled on the virtual drive.

This virtual drive becomes a snapshot repository. Use it only for storing snapshot-related data.

4. If you click **No**, the snapshots are not enabled on the selected virtual drive.

CAUTION: After you enable the snapshots on this virtual drive, you cannot change the allocated percentage of capacity or the snapshot repository without first disabling the snapshots and losing any snapshot data.

11.2.17 Snapshot Base Details

You can view the snapshot base details of the virtual base drive.

Perform the following steps to view the details of the snapshot of the virtual base drive:

1. Select the **Logical** tab on the MegaRAID Storage Manager window.
2. Click a base virtual drive in the left frame.

After you select the base virtual drive, the base virtual drive information appears in the right frame of the **Properties** tab under **Snapshot Base Details** (marked in pink in the following figure).

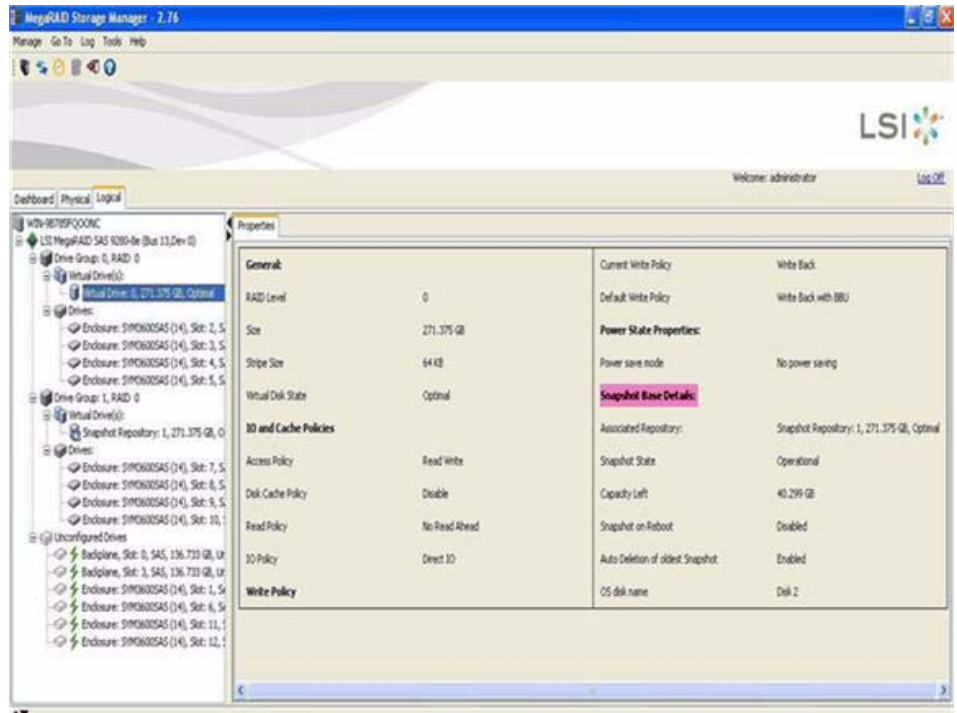


Figure 292: Snapshot Base Details

11.2.18 Manage Snapshots

You can create snapshots, delete snapshots, create views, and also edit, pause, or delete schedules using the Manage Snapshots wizard.

1. Select **Go To>>Virtual Drive>>Manage Snapshot wizard** on the menu bar, as shown in the following figure.

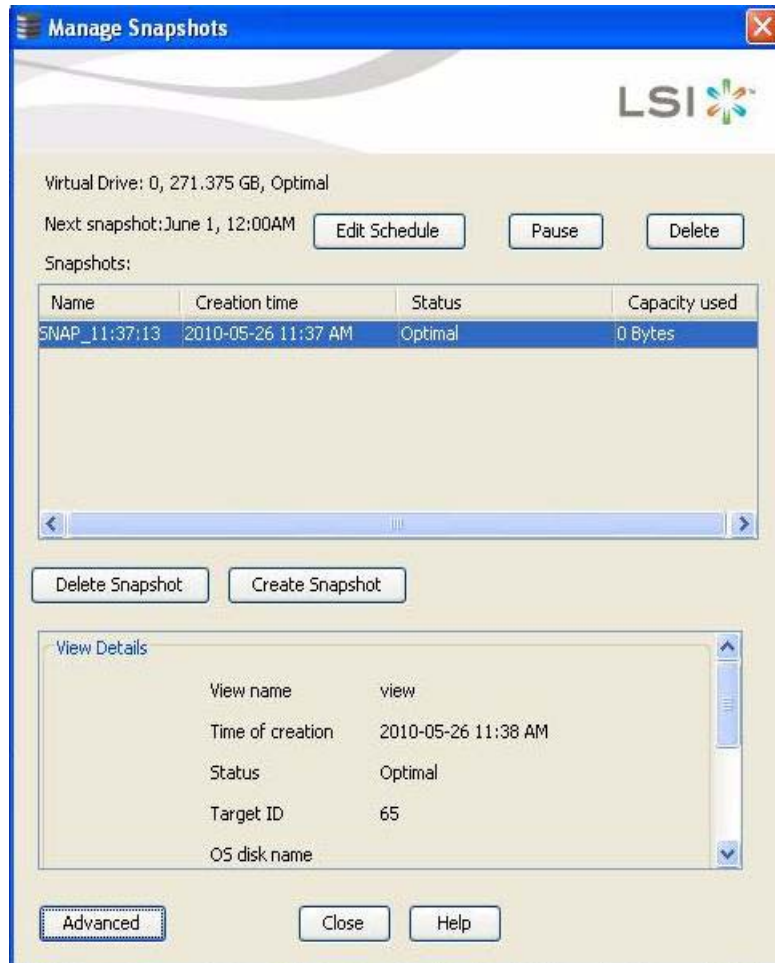


Figure 293: Manage Snapshots Dialog

You can edit the snapshot schedule using the **Edit Schedule** button, pause the snapshot schedule using the **Pause** button, and delete the snapshot schedule using the **Delete** button.

The **Snapshots** table displays the snapshot details.

- The **Name** column displays the name of the snapshot.
- The **Status** column displays the status of the snapshot.
- The **Capacity Used** column displays the capacity consumed by the snapshot.

You can create the snapshot by clicking **Create Snapshot**, and delete the snapshot by clicking **Delete Snapshot**.

In the **View Details** frame, you can create a view by clicking **Create View** and edit the settings for automatic snapshots by clicking **Advanced**.

11.2.19 Editing Schedule

You can edit the schedule using the **Edit Schedule** dialog. You can change the frequency of the snapshot, the day in which the snapshot needs to be taken, and the time during which the snapshot needs to be taken.

1. Click the **Edit Schedule** button in the Manage Snapshots wizard in [Figure 293](#), if you want to edit the snapshot schedule.

The **Edit Schedule** dialog appears, as shown in the following figure.

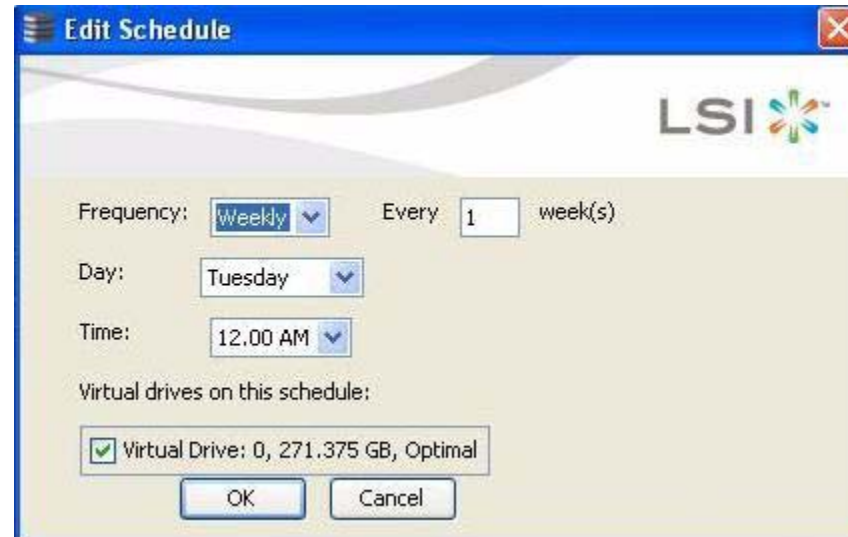


Figure 294: Edit Schedule Dialog

- In the **Frequency** field, use the drop-down list to edit the frequency of the snapshot already taken. The frequency can be daily or weekly.
- In the **Day** field, use the drop-down list to edit the day of the snapshot already taken. The days can be from Monday through Sunday.
- In the **Time** field, use the drop-down list to edit the time of the snapshot already taken.

After you select all of the above fields, the virtual drives matching these fields appear in the **Virtual drives on this schedule** check box.

2. Select the **Virtual drives on this schedule** check box, and click **OK**.

The virtual drive details are edited.

11.2.20 Advanced Settings

You can edit the settings for the automatic snapshots. You can automatically delete the oldest snapshot, or stop taking snapshots, or take a snapshot on every reboot.

1. Click **Advanced** in the Manage Snapshots wizard as shown in [Figure 293](#).

The Advanced dialog appears, as shown in the following figure.



Figure 295: Advanced Dialog

You can edit the settings by selecting one of the following options:

- The **Automatically delete the oldest snapshot** option, if you want to delete the oldest snapshot.
- The **Stop taking snapshots** option, if you want to stop taking snapshots.
- The **Take a snapshot on every reboot (recommended for boot virtual drive)** check box, if you want a snapshot on every reboot.

2. Click **OK**.

The settings are edited.

11.2.21 Create View Using Manage Snapshots Wizard

You can create views using the **Create View** button present in the **Manage Snapshots** area under the **View Details** field. The view provides the snapshot details of the virtual drive available at that particular time.

Follow these steps to create views of the snapshots.

1. Click **Create View** button in the Manage Snapshots wizard, as shown in [Figure 293](#).

The **Create View** dialog appears, as shown in the following figure.



Figure 296: Create View Dialog

2. In the **View name** field, enter the view name. For example, `view`.
3. In the **Write capacity** field, use the drop-down list to allocate capacity for writes to the view.
4. Click **OK**.

The capacity is allocated for writes to the view.

11.2.22 Viewing Snapshot Details

If the view details of the snapshot are available at that particular time for the virtual drive, these details appear under the view details in the Manage Snapshots wizard, as shown in the following figure.

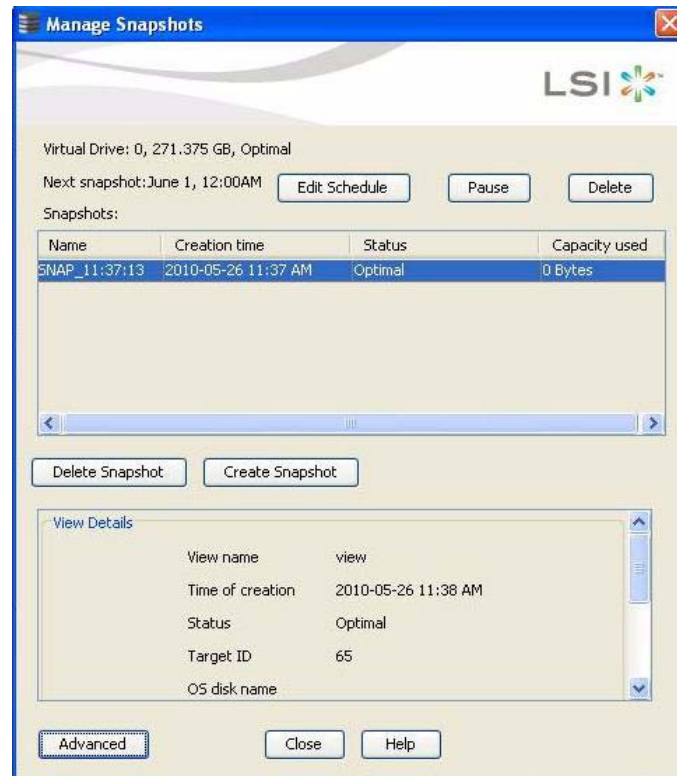


Figure 297: Manage Snapshots Dialog

11.2.23 No View Details for Snapshot

When no view for the snapshot exists, the following message appears in the **View details** area, as shown in the following figure.



Figure 298: Manage Snapshots Dialog - No View Present for the Snapshot

11.2.24 No Snapshot Schedule

When there are virtual drives with no snapshot schedule, the following message appears, as shown in Figure 299.

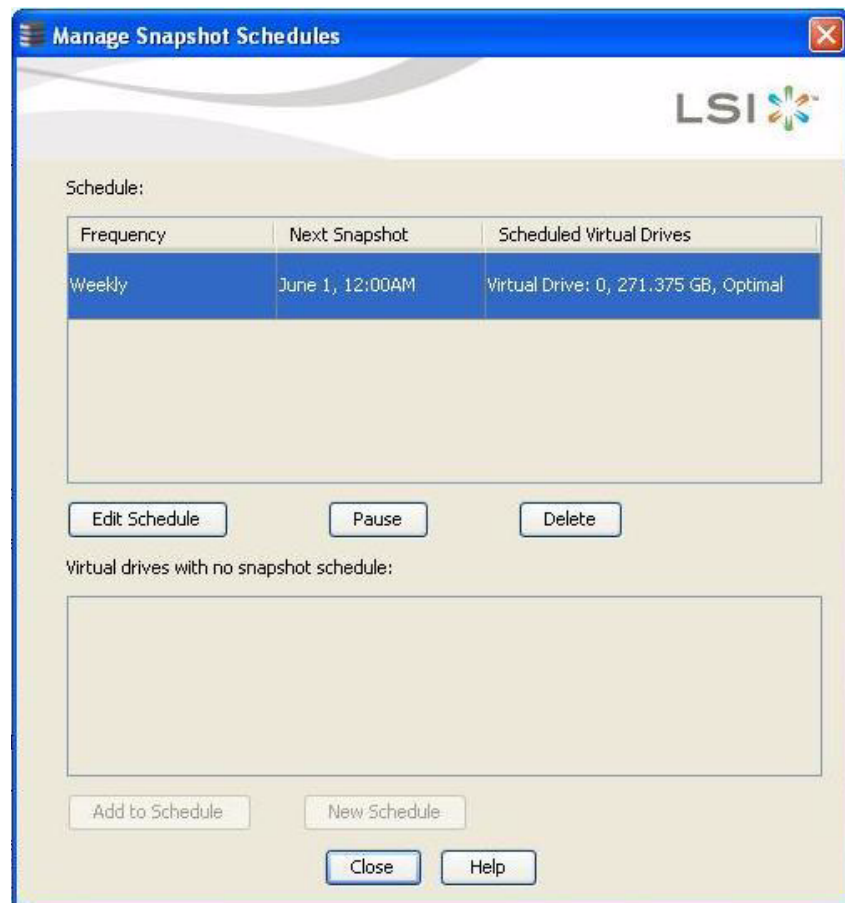


Figure 299: Manage Snapshot Schedules Dialog - Virtual Drives with No Snapshot Schedule

- Click **Add to Schedule** to add a snapshot schedule.
- Click **New Schedule** to add a new snapshot schedule.

11.2.25 Graphical Representation of Repository Virtual Drive

To view the graphical representation of the repository virtual drive, perform the following steps:

1. Click the Logical view on the main menu window.
2. Click the Snapshot Repository virtual drive in the left frame.
3. Click the Snapshots tab in the right frame. The following figure appears, which shows the graphical representation of the virtual drive details.

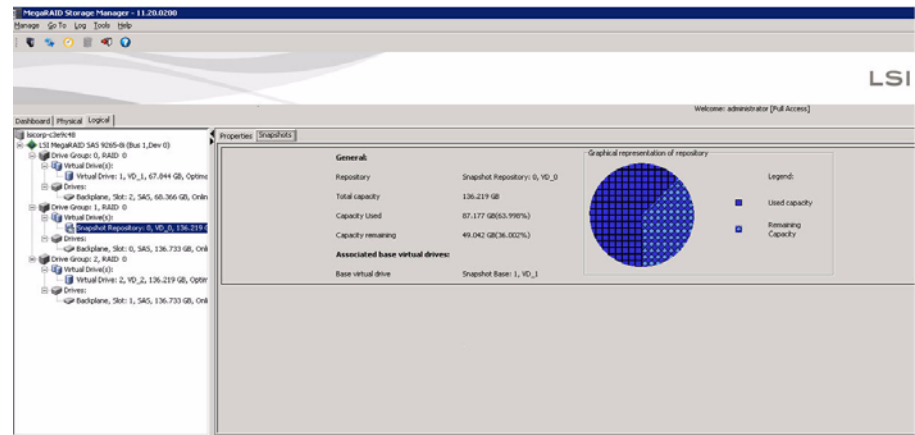


Figure 300: Repository Virtual Drive Details

11.2.26 Deleting a Snapshot

NOTE: You can delete only the oldest snapshot.

Follow these steps to delete a snapshot.

1. Click the **Logical** tab on the main menu window in the Logical view.
2. Select a Required Base virtual drive from the list of virtual drives in the left frame.
3. Select **Go To >> Virtual Drive >> Manage Snapshots** on the menu bar. The window that appears shows the Snapshot Base details and any existing snapshots.
4. Click the oldest snapshot in the timeline.
5. Click the **Delete Snapshot** button.

This action deletes the oldest snapshot.

11.3 Disabling MegaRAID Recovery

Follow these steps to disable MegaRAID recovery.

1. Click the Logical tab on the main menu window in the Logical view.
2. Select and highlight a required base virtual drive from the list of virtual drives.
3. Select **Go To >> Virtual Drive >> Disable MegaRAID Recovery** on the menu bar. The following confirmation dialog appears.

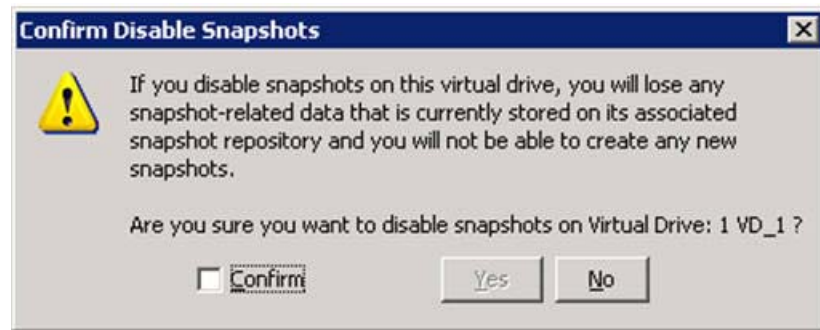


Figure 301: Confirm Disable Snapshots

4. Select the **Confirm** check box if you want to disable snapshots. When you select this check box, the Yes button gets enabled. The snapshots get disabled on the virtual drive. If you click No, the snapshots will not be disabled on the selected virtual drive.

11.4 CacheCade Advanced Software

MegaRAID CacheCade improves application performance by expanding the MegaRAID read-caching capacity and introducing write caching. The CacheCade feature uses high-performing CacheCade software as a secondary tier of cache to provide faster reads and to maximize transactional I/O performance.

Using the CacheCade software as controller cache allows for very large data sets to be present in cache, delivering up to a performance improvement that is 50 times greater than regular cache in read-intensive applications, such as online transaction processing (OLTP), and file and Web server workloads. The solution is designed to accelerate the I/O performance of HDD-based drive groups while only requiring a small investment in CacheCade software technology.

To support full-throughput for multiple direct-attached CacheCade software, this feature reduces IO-processing overhead in the 2108-chip-based MegaRAID controllers. CacheCade offers performance equivalent to flash-based controllers and better performance for RAID 5 and RAID 6 when compared to Fusion I/O.

11.4.1 Using the CacheCade 2.0 SSD Read Caching Advanced Software

CacheCade 2.0 provides you with read caching capability

Perform the following steps to use the CacheCade advanced software.

1. Click a RAID controller icon in the left frame.
2. Select **Controller >> Create CacheCade** on the menu bar, as shown in [Figure 302](#).

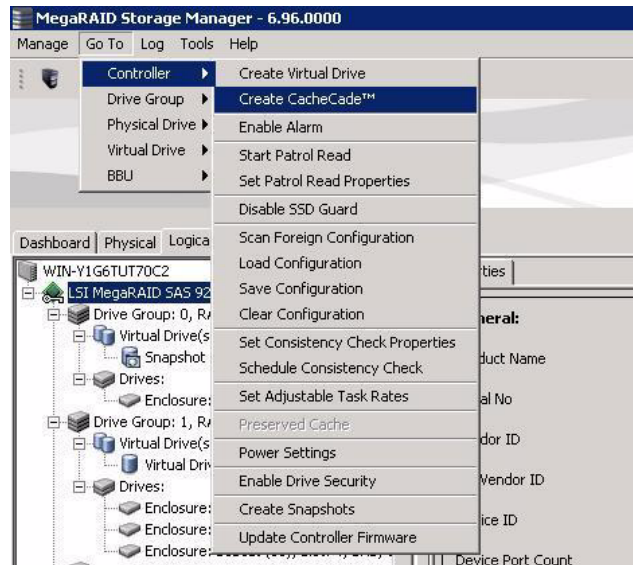


Figure 302: Create CacheCade Menu Option

The wizard dialog appears.

3. Click on unconfigured CacheCade - SSD Caching drives in the left frame to select the drives for the CacheCade drive group, as shown in [Figure 303](#).

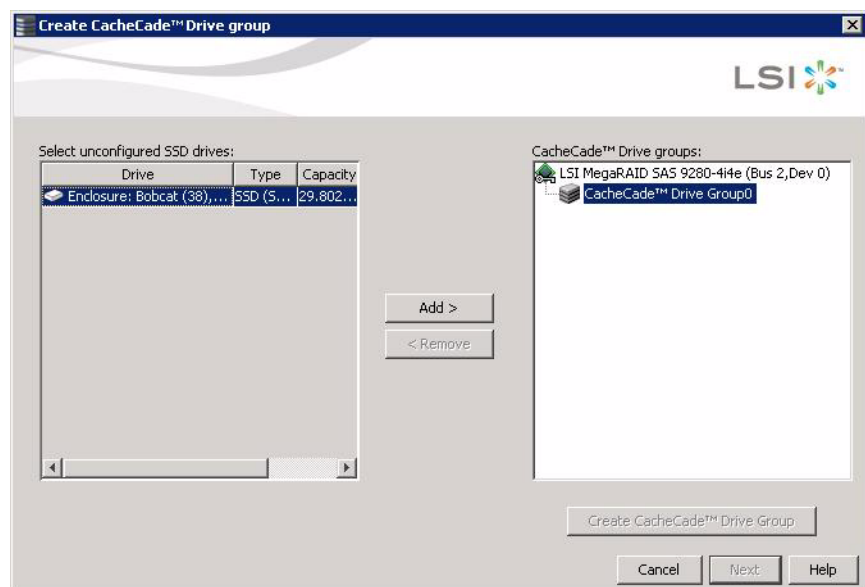


Figure 303: Create CacheCade™ Drive group Dialog

After you select the unconfigured drives, the **Add >** button is available.

- Click **Add >** to move the selected drives to the drive group in the right frame, as shown in [Figure 304](#).

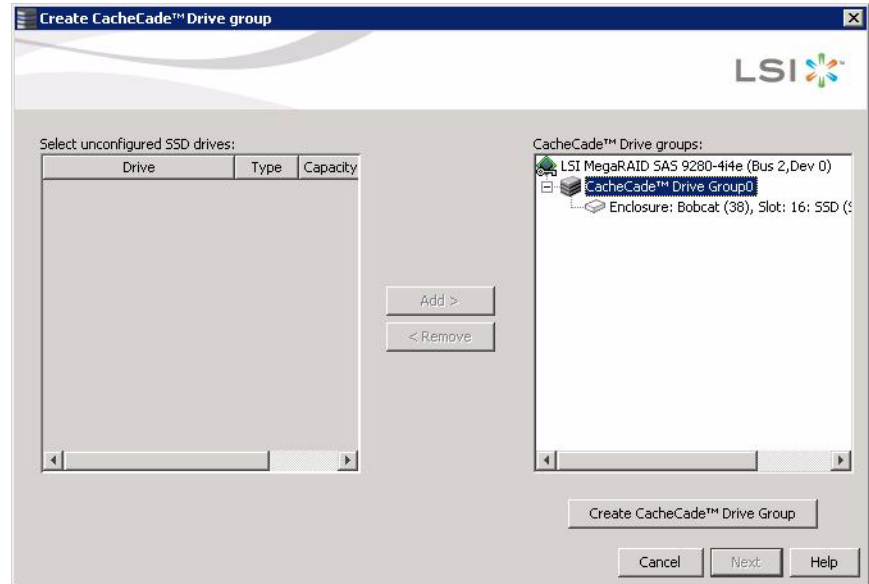


Figure 304: CacheCade Drive Group Dialog

After you move the selected drives, the **Create CacheCade™ Drive Group** button is available.

- Click **Create CacheCade™ Drive Group**.
- Click **Next**.

Use the next dialog that appears to select parameters for the cache disk.

- Enter a name for the CacheCade - SSD Caching virtual drive in the **CacheCade virtual drive name** field, and click **Create**.

Depending on the number of drives, you might have the option to set the capacity of the CacheCade - SSD Caching drive.

The CacheCade drive group icon appears in the menu dialog, as shown in [Figure 305](#).

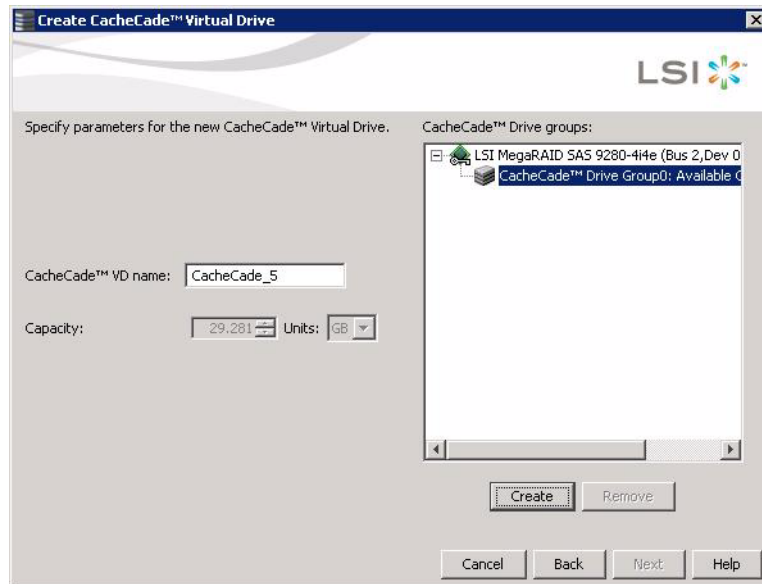


Figure 305: Create CacheCade™ - Summary Dialog

8. Click **Next**.

The summary dialog appears, as shown in [Figure 306](#). This dialog displays the drive group name, the number of drives, the total capacity, the free capacity, the CacheCade virtual drive name, and the capacity being used.

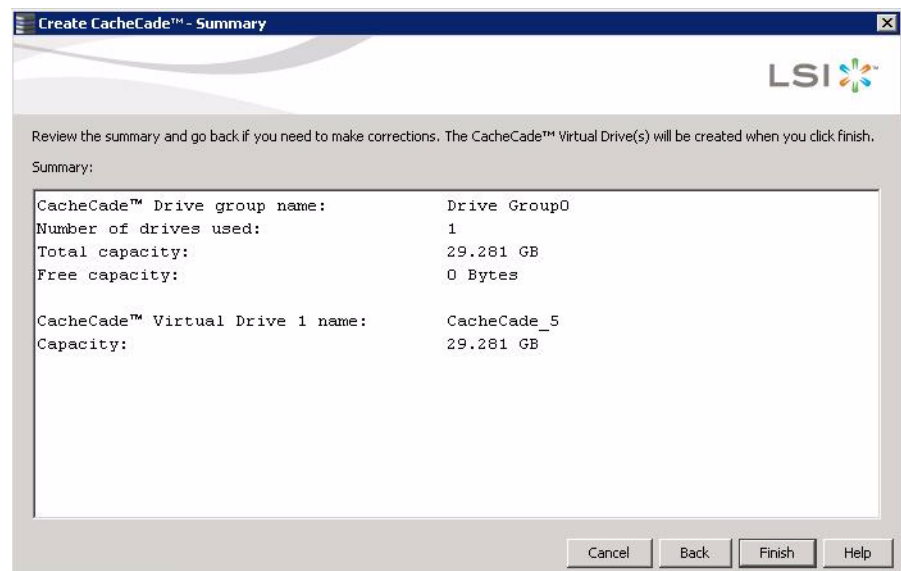


Figure 306: CacheCade Virtual Drive Summary Dialog

9. Click **Finish**.

A confirmation message displays after the CacheCade virtual drive is successfully created.

The CacheCade drive icon appears next to the RAID controller in the left frame, as shown in the following figure.

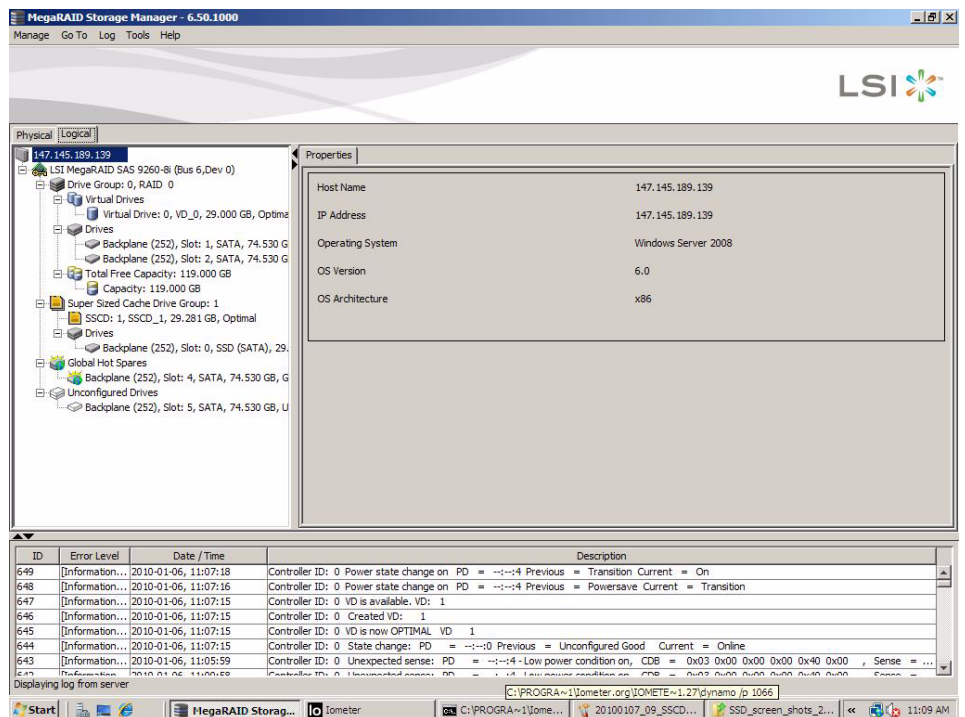


Figure 307: CacheCade Drive Icon in MegaRAID Storage Manager

11.4.2 Using the CacheCade Pro 2.0 SSD Read/Write Caching Software

CacheCade Pro 2.0 SSD Read/Write Caching software provides you with read and write caching capability.

NOTE: The MegaRAID firmware has the provision to monitor IO performance; changes have been made to accommodate the CacheCade Pro 2.0 software statistics. The CacheCade Pro 2.0 software metrics are captured for each logical drive that has CacheCade enabled. The CacheCade Pro 2.0 software gathers information about the cache windows allocated for a logical drive, the number of new windows allocated in this metrics collection period, the number of windows that are actively used, and the window hit rates.

Perform the following steps to use the CacheCade Pro 2.0 SSD Read/Write Caching software:

1. Right click on a controller in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create CacheCade SSD Caching**.

OR

Select a controller and select **Go To>>Controller>>Create CacheCade SSD Caching** in the menu bar. The CacheCade SSD Caching Wizard appears, as shown in the following figure.

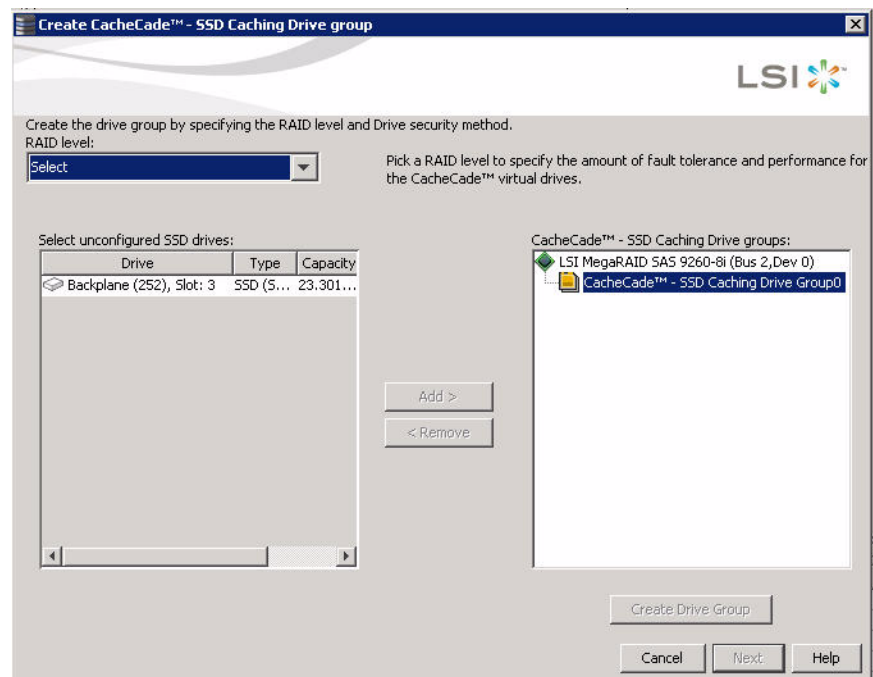


Figure 308: CacheCade SSD Caching Wizard - First Screen

2. Select a RAID level for the CacheCade virtual drive in the **RAID level** field.
3. Select an unconfigured SSD drive, for the selected RAID level, from **Select unconfigured SSD Drives** in the left frame.

After you select an unconfigured SSD Drive, the **Add** button is enabled.

4. Click **Add** to add the selected drive to the CacheCade - SSD Caching Drive groups in the right frame.

After you click **Add**, the **Create Drive Group** button is enabled.

5. Click **Create Drive Group**. The newly created drive group appears in CacheCade SSD Caching Drive groups in the right frame.
6. Click **Next**. The next wizard screen appears.

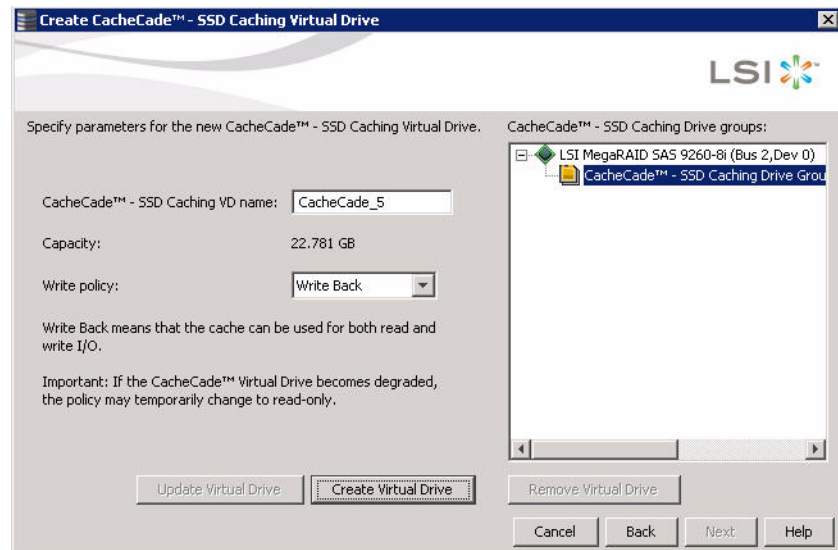


Figure 309: Parameters for CacheCade SSD Caching Virtual Drive

7. Enter a name for the CacheCade virtual drive in the **CacheCade - SSD caching VD name** field.
8. Select a write policy from the **Write policy** drop-down list. A description of the selected write policy appears below.
9. Click **Create Virtual Drive**. The newly created virtual drive appears in the CacheCade SSD Caching Drive groups in the right frame. The **Remove Virtual Drive button** is enabled. You can select the newly created virtual drive and click **Remove Virtual Drive** to delete the virtual drive.
10. Click **Next**. The summary screen appears.

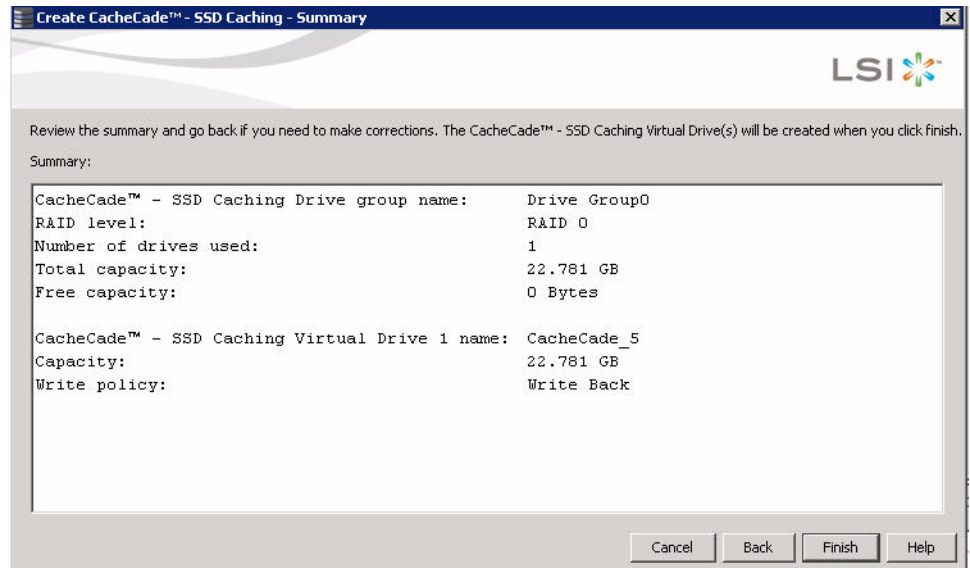


Figure 310: Create CacheCade - SSD Caching - Summary

This screen displays the drive group name, the RAID level, the number of drives, the total capacity, the free capacity, the CacheCade virtual drive name, the capacity being used, and the write policy.

11. Click **Finish**.

A confirmation message displays after the CacheCade virtual drive is successfully created. The CacheCade drive icon appears next to the RAID controller in the left frame in the MegaRAID Storage Manager window.

11.4.2.1 Modifying the CacheCade Virtual Drive Properties

You can modify the name and the write policy of a CacheCade virtual drive any time after a CacheCade virtual drive is created. Perform the following steps to change the virtual drive properties:

1. Right-click on a controller in the device tree in the left frame of the MegaRAID Storage Manager window and select **Set Virtual Drive Properties**.

OR

Select a controller and select **Go To>>Virtual Drive>>Set Virtual Drive Properties**. The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.

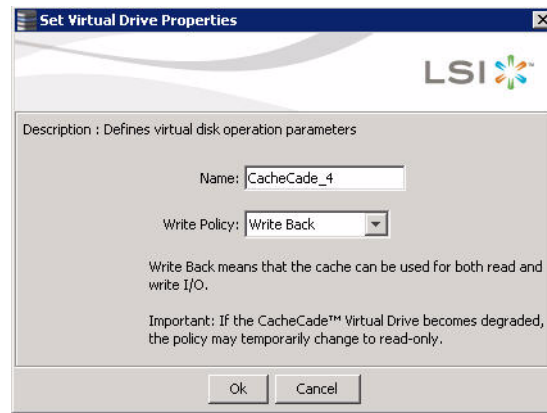


Figure 311: Set Virtual Drive Properties

2. Edit the name of a CacheCade virtual drive in the **Name** field.
3. Select a write policy from the **Write Policy** drop down list.
4. Click **OK**. A confirmation dialog appears with a warning note. Select the **Confirm** check box and click **OK**.

11.4.2.2 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade SSD Caching virtual drive. This option is only available when the virtual drive's caching is currently disabled.

Perform the following steps to enable SSD caching on a virtual drive:

1. Right click on a virtual drive in the left frame of the MegaRAID Storage Manager window and select **Enable SSD Caching**.

OR

Select a virtual drive and select **Go To>>Virtual Drive>>Enable SSD Caching**. The **Enable SSD Caching** dialog appears, as shown in the following figure.




Figure 312: Enable SSD Caching

2. Click **OK** to enable caching for that virtual drive.

11.4.2.3 Disabling SSD Caching on a Virtual Drive

You can disable caching on a virtual drive. When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade SSD Caching virtual drive is removed. This option is only available when the virtual drive's caching is currently enabled.

Perform the following steps to enable SSD Caching on a virtual drive:

1. Right click a virtual drive in the left frame of the MegaRAID Storage Manager window and select **Disable SSD Caching**.

OR

Select a virtual drive and select **Go To>>Virtual Drive>>Disable SSD Caching**. The **Disable SSD Caching** dialog appears, as shown in the following figure.




Figure 313: Disable SSD Caching

2. Select the **Confirm** check box and click **OK** to disable caching for that virtual drive.

11.4.2.4 Enabling/Disabling SSD Caching on Multiple Virtual Drives

You can enable or disable SSD caching on multiple virtual drives at one go.

Perform the follow steps to enable or disable SSD caching on multiple drives:

1. Right click a controller in the left frame of the MegaRAID Storage Manager window and select **Manage SSD Caching**.

OR

Select a controller and select **Go To>>Controller>>Manage SSD Caching**. The **Manage SSD Caching** dialog appears, as shown in the following figure.

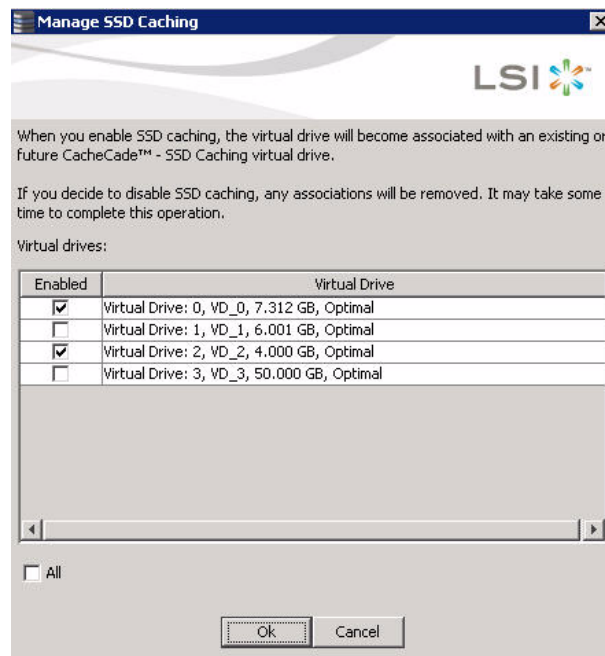


Figure 314: Manage SSD Caching

The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.

2. Select or deselect a check box to change the current setting of a virtual drive.
3. Click **OK**.

If you select the **All** check box, all the virtual drives are enabled. If you deselect the **All** check box, all the virtual drives are disabled.

If you disable SSD caching on a virtual drive, the **Disable SSD Caching** dialog appears.

4. Select the **Confirm** check box and click **OK** to enable/disable SSD caching on the selected virtual drives.

11.4.2.5 Modifying a CacheCade Drive Group

To modify an existing CacheCade SSD caching drive group, you need to first delete the drive group and then create a new CacheCade drive group.

11.4.2.6 Clearing Configuration on CacheCade Pro 2.0 Virtual Drives

You can clear all existing configurations on a selected controller that has CacheCade Pro 2.0 virtual drives.

1. Right click on a controller in the left frame of the MegaRAID Storage Manager window and select **Clear Configuration**.

OR

Select a controller and select **Go To>>Controller>>Clear Configuration**. The Confirm Clear Configuration dialog appears as shown, in the following figure.

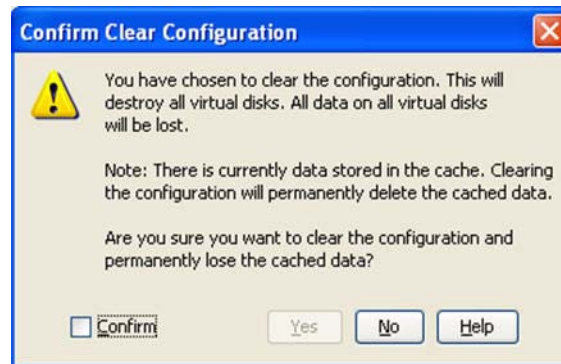


Figure 315: Confirm Clear Configuration

2. Select the **Confirm** check box and click **Yes**.
3. If the cache becomes inconsistent before the clear configuration operation is performed, the firmware returns an error code. The **Confirm Loss of Cache** dialog appears as a follow up dialog to the **Confirm Clear Configuration** dialog.
4. Select the **Confirm** check box and click **Yes**.

11.4.2.7 Removing Blocked Access

At times, an error may occur in the CacheCade virtual drive and this causes a blocked access to the associated virtual drive.

An icon appears in front of the affected virtual drive, next to the *Optimal* status.

It is advisable to wait for sometime for the error in the CacheCade virtual drive to get sorted. You can also try to solve the error in the CacheCade virtual drive and bring it back to an optimal status. Once the CacheCade virtual drive is in an optimal status, the blocked virtual drive returns to its former access policy automatically.

If it is not possible to bring the CacheCade virtual drive to its optimal status, follow these steps to remove the blocked access from the virtual drive:

1. Right click on the icon on the virtual drive with the blocked access and select **Remove Blocked Access**. The **Confirm Remove Blocked Access** dialog appears, as shown in the following figure.



Figure 316: Confirm Remove Blocked Access

2. Select the **Confirm** check box and click **Yes**.

11.4.2.8 Deleting a Virtual Drive With SSD Caching enabled

You can delete a virtual drive that has SSD caching enabled on it.

Perform the following steps to delete the virtual drive:

1. Right click on a CacheCade virtual drive and select **Delete Virtual Drive**.

OR

Select a CacheCade virtual drive and click **Go To>>Virtual Drive>>Delete Virtual Drive**. The Confirm Delete Virtual Disk dialog appears, as shown in the following figure.



Figure 317: Confirm Delete Virtual Disk

2. Select the **Confirm** check box and click **Yes**.

NOTE: If you select the **Force the delete to complete quickly** check box to delete the virtual drive, the data is not flushed before deleting the virtual drive. In this scenario, if you create this virtual drive after deleting it, there will be no data available.

11.5 FastPath Advanced Software

MegaRAID FastPath is a high-performance IO accelerator for the CacheCade 2.0 SSD Read Caching software drive groups connected to a MegaRAID controller card. CacheCade2.0 SSD Read Caching software has a read performance advantage over HDDs and uses less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 6Gb/s MegaRAID SATA+SAS controller.

The FastPath feature supports full optimization of the CacheCade 2.0 SSD Read Caching software and hard disk drive (HDD) virtual disk groups to deliver an improvement in read and write IOPS that is three times greater than MegaRAID controllers not using FastPath technology. Also, the FastPath advanced software is faster and more cost-effective than current flash-based adapter card solutions.

11.5.1 Setting FastPath Options

Perform the following steps to use the FastPath advanced software:

1. Select the **Logical** tab on the MegaRAID Storage Manager window for the Logical view.
2. Select a virtual drive icon in the left frame.
3. Select **Virtual Drive >> Set Virtual Drive Properties** on the menu bar, as shown in the following figure.

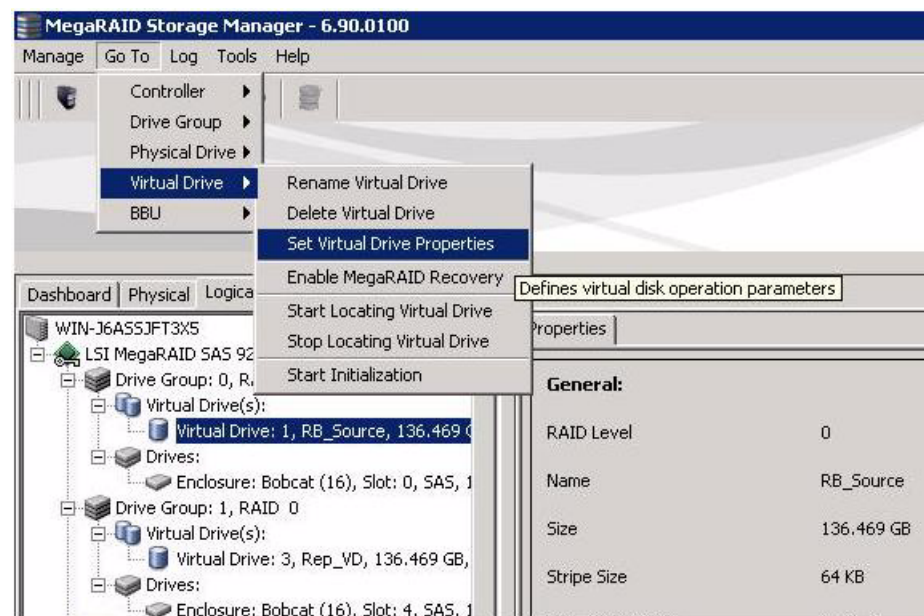


Figure 318: Set Virtual Drive Properties Menu

The Set Virtual Drive Properties dialog appears. It shows the default settings for the FastPath advanced software:

- Write Policy: Write Thru

- IO Policy: Direct IO
 - Read Policy: No Read Ahead
 - Dish Cache Policy: Enabled
 - Strip Size: 64KB
4. Click **OK**.
A confirmation dialog displays.
 5. Select the **Confirm** check box, and click **Yes** to confirm that you want to set the virtual drive properties.

11.6 LSI SafeStore Encryption Services

LSI SafeStore Encryption Services offer the ability to encrypt data on the drives and use the drive-based key management to provide data security. This solution provides data protection in the event of theft or loss of physical drives. If you remove a self-encrypting drive from its storage system or the server in which it resides, the data on that drive is encrypted, and becomes useless to anyone who attempts to access it without the appropriate security authorization.

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration using the SafeStore Encryption Services advanced software.

The SafeStore Encryption Services advanced software provides drive security to create secure virtual drives by using **External Key Management (EKM)** and **Local Key Management (LKM)**.

11.6.1 Enabling Drive Security using EKM

EKM is used for key management when large number of systems are deployed. You can automate and manage the life cycle of keys and unlock configurations using EKM.

Another important feature of EKM is that you can use it without human intervention to perform operations like drive migration and controller replacement.

MegaRAID accomplishes the task of obtaining keys by interacting with the EKM agent. The EKM agent talks to the EKM server (EKMS) through a network and gets the security key for the controller.

Keys are retrieved or created to perform the following tasks:

- Create secure Virtual drives.
- Insert drives to replace failed drives in a secure configuration.
- Re-key the system based on EKMS policies or user request.
- Gain access to a secured configuration during boot.
- Unlock and import secured drives during migration.

Perform the following configurations to enable the drive security to create secure virtual drive using the EKM mode with the support of EKM servers.

- EKM mode is supported by MegaRAID Storage Manager, and EKMS is present.
- EKM mode is supported by MegaRAID Storage Manager, and EKMS is not present.
- Change the current security settings, or switch between the modes.
- Change the security settings when the user is in EKM, and wants to switch to LKM.
- Import Foreign Drives.

11.6.2 Supporting EKM Mode

When you choose EKM for drive security, and decide to configure, EKM mode is supported, and EKMS is present, the application responds to different behaviors depending on the scenarios at that particular time.

The first scenario occurs when EKM is enabled, and the second scenario occurs when EKM is enabled and EKMS is present. The details of this scenario are described further in this section.

Perform the following steps to configure, EKM mode is supported, and EKMS is present:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Choose any one of the following options to arrive at the Drive Security Choose Mode wizard.
 - Select **Go To >> Controller >> Enable Drive Security** in the main navigation bar in the top portion of the MegaRAID Storage Manager window.
 - Right-click the controller icon, and click **Enable Drive Security** menu.

The Drive Security Choose Mode wizard appears, as shown in the following figure.

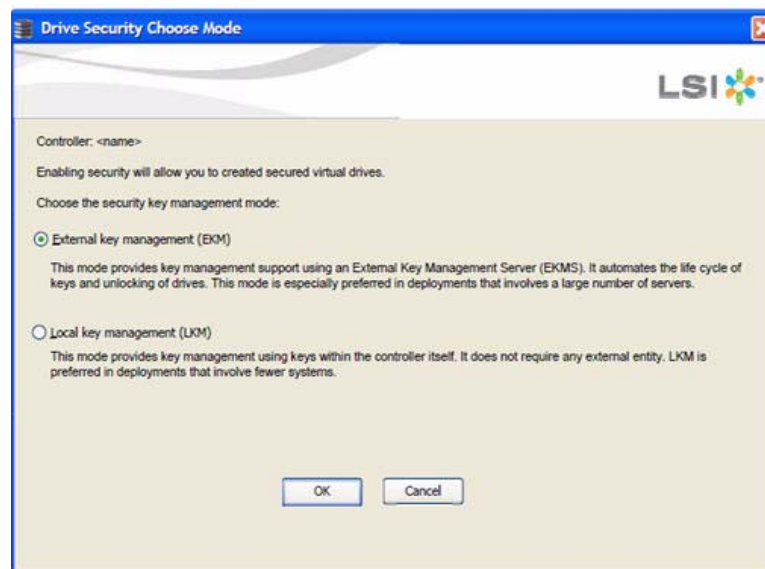


Figure 319: Drive Security Choose Mode Dialog

3. Select the **External Key management (EKM)** radio button.
4. Click **OK**.

After you click **OK**, two scenarios occur, based on the availability of EKMS.

- Scenario # 1 – **EKM is enabled**

When EKM is enabled, you can see the confirmation message as shown in the following figure. The message displayed makes sure that the security is enabled on the controller using EKMS.



Figure 320: Confirm External Key Management Mode Enabled

NOTE: On the right frame of the controller properties, in **Drive Security Properties**, you can see the key management mode is EKM.

- Scenario # 2 – **EKM is selected, and EKMS is not present**

When EKM is selected and EKMS is not present, you must restart the system. When you restart the system, the system restart message appears as shown in the following figure.




Figure 321: System Restart Required Message

When the system restarts, the boot agent generates the security keys for the controller.

NOTE: If the MegaRAID Storage Manager application does not support EKM, the EKM option is greyed out.

11.6.3 Change Security Settings- LKM

When you select the **Change Security Settings** option in the **Change Security wizard** as shown in [Figure 322](#), two options are provided for you. If you select the first option, **Change current security settings**, you can change the drive security settings on the controller. If you select the second option, **Switch to External Key Management (EKM) mode**, you can switch from LKM mode to EKM mode.

Perform the following steps to configure the Change Security Settings wizard.

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Choose any one of the following options to enable the **Change Security Settings** wizard.
 - Select **Go To >> Controller >> Change Security Settings** in the main navigation bar present in the top portion of the MegaRAID Storage Manager dialog.
 - Right-click the controller icon, and click **Change Security Settings**.

The **Change Security Settings** wizard appears as shown in the following figure.

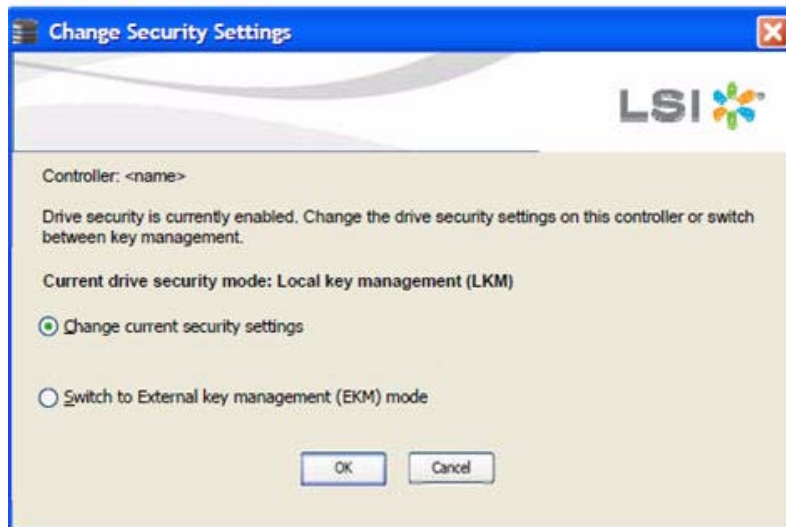


Figure 322: Change Security Settings Dialog

3. Select the **Change current security settings** radio button, if you want to change the drive security settings on the controller using the LKM mode.

When you select this option, the details of the change current security settings given in [Section 11.6.10, Changing the Drive Security Settings](#) is displayed.

4. Select the **Switch to External key management (EKM) mode** radio button, if you want to switch over from LKM mode to EKM mode.
5. Click **OK**.

After you click **OK**, the Authentication Drive Security Key dialog appears.

11.6.3.1 Authentication Drive Security Key

In the Authentication Drive Security Key dialog, enter the current LKM authentication drive security key in the **Current Security Key** field.



Figure 323: Authentication Drive Security Key

If you enter a valid authentication key, a confirmation message indicating that EKM is enabled is displayed as shown in the following figure.



Figure 324: Confirm External Key Management (EKM) Mode Enabled Message

Based on the application's response at that particular time, you come across two scenarios. Scenario #1, if **EKMS is selected and EKMS is not present**. Scenario #2, if you enter a wrong security key in the **Authentication Drive Security Key** dialog box in [Figure 323](#).

■ Scenario # 1 – **EKM is selected and EKMS is not present**

When EKM is selected and EKMS is not present, restart the system, and during re-boot the boot agent generates the security keys.

When you restart the system, the system restart message appears, as shown in the following figure.

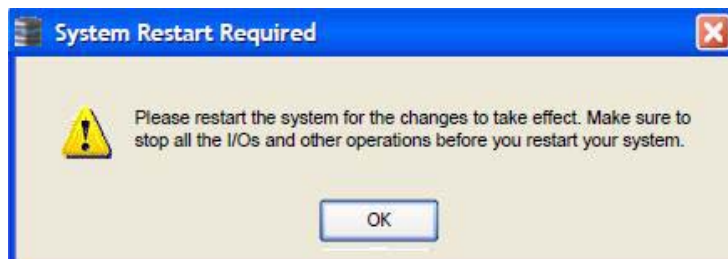


Figure 325: System Restart

■ Scenario #2 – Entering a wrong security key

If you enter a wrong security key in the **Authentication Drive Security Key** dialog box in [Figure 323](#), an error message is displayed, as shown in the following figure.




Figure 326: Error Message When the User Enters Wrong Security Key

11.6.4 Change Security Settings - EKM

When you select the **Change Security Settings** in the Change Security wizard as shown in [Figure 327](#), two options are offered to you.

If you select the first option, **Change current security settings**, you can traverse to the EKM mode.

If you select the second option, **Switch to Local Key Management (LKM)** mode, you can switch over to LKM mode.

Perform the following steps to configure the Change Security Settings wizard:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Choose any one of the following options to enable the Change Security Settings wizard.
 - Select **Go To >> Controller >> Change Security Settings** in the menu bar present in the top portion of the MegaRAID Storage Manager dialog.
 - Right-click the controller icon, and select **Change Security Settings**.

The Change Security Settings wizard appears, as shown in the following figure.



Figure 327: Change Security Settings Dialog

When you select the **Change current security settings** radio button, a different EKM key for the controller is generated.

Based on the application's response two scenarios occur, Scenario #1, when **EKM server is present**, Scenario #2, when **EKM server is not present**

■ Scenario # 1 – **EKM server is present**

If the EKM server is present, a confirmation message indicating that the drive security is enabled on the controller as shown in the following figure is displayed.



Figure 328: Confirm External Key Management (EKM) Mode Enabled Message

■ Scenario # 2 – **EKM server is not present**

If the EKM server is not present, restart your system, and during re-boot, the boot agent generates the security keys.

When you restart the system, the system restart message appears, as shown in the following figure.

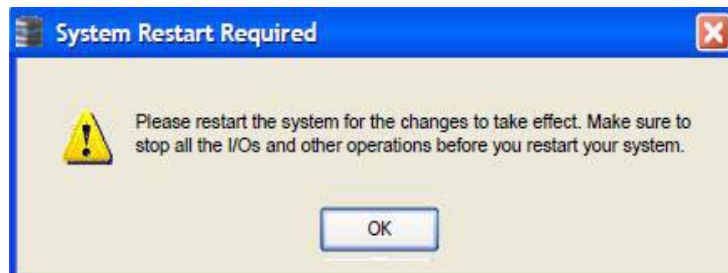


Figure 329: System Restart Required Message

NOTE: If you select the **Switch to Local key Management (LKM)** mode in the Change Security Settings Wizard in [Figure 327](#), you can view the details of LKM in [Enabling Drive Security using LKM](#).

11.6.5 Importing Foreign Drives

You can import the foreign drives in the LKM or EKM mode. Depending on the mode selected and the secured drives, the following scenarios occur.

- Import Foreign Drives – **LKM mode, EKM Secured Locked Drives are present.**
- Import Foreign Drives – **EKM mode, LKM secured locked drives are present.**
- Import Foreign Drives – **EKM mode, EKM secured locked drives are present.**

Perform the following steps to configure the Scan Foreign Configuration wizard:

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

2. Choose any one of the following options to enable the Scan Foreign Configuration wizard.
 - Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window, and select **Scan Foreign Configuration**.
 - Select **Go To >> Controller >> Scan Foreign Configuration** in the main navigation bar present in the top portion of the MegaRAID Storage Manager window.

If a foreign configuration is detected, the “Foreign Configuration Detected” dialog appears. If you select **Import** and click **OK**, the Unlock Foreign Drives wizard appears, as shown in [Figure 331](#).

11.6.6 Importing Foreign Drives to LKM

When the virtual drive is secured with the EKM mode, and if you change the security settings from the EKM mode to LKM mode, and want to import those foreign configured virtual drives, the following dialog appears.



Figure 330: Importing Foreign Drives- EKM to LKM

11.6.7 Importing Foreign Drives to EKM

When the virtual drive is secured with the LKM mode, and you want to switch from LKM to EKM, and if the virtual drive is in Foreign state, and if you want to import those foreign configuration, you must provide the LKM key for each virtual drive.



Figure 331: Unlock Foreign Drives

- The **Locked drives** table displays the drives that are locked with the LKM mode.
- The **Security Key** field displays the LKM security key for the virtual drive.

After you enter the security key in the **Security Key** field, the system tries to unlock each of the locked foreign drives using the security key. If at least one drive is unlocked, then the Import preview window appears. If no drives are unlocked, the Incorrect security Key dialog appears.

11.6.8 Importing Foreign Drives to EKM

When a virtual drive is secured with EKM mode, and if that virtual drive is in Foreign state, and if you switch to EKM mode and if you want to import to the foreign configuration, then the dialog appears, as shown in the following figure.



Figure 332: Switch to EKM Mode Dialog

11.6.9 Enabling Drive Security using LKM

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration using the SafeStore Encryption Services advanced software.

To enable security on the drives, you need to perform the following actions to set drive security:

- Enter a security key identifier.

A security key identifier appears whenever you have to enter a security key. If you have more than one security key, the identifier helps you determine which security key to enter.

- Enter a security key.

After you create a security key, you have the option to create secure virtual drives using the key. You have to use the security key to perform certain operations.

You can improve security by entering a password. To provide additional security, you can require the password whenever anyone boots the server.

Perform the following steps to enable drive security.

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Select **Go To >> Controller >> Change Drive Security >> Enable**.

The Enable Drive Security – Enter Security Key Details dialog appears, as shown in the following figure.

Figure 333: Enable Drive Security - Enter Security Key Details Window

3. Either use the default security key identifier, or enter a new security key identifier.

NOTE: If you create more than one security key, make sure that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

4. Either click **Suggest Security Key** to have the system create a security key, or you can enter a new security key.
5. Enter the new security key again to confirm, as shown in the following figure.

CAUTION: If you forget the security key, you will lose access to your data. Be sure to record your security key information. You might need to enter the security key to perform certain operations.

The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.

NOTE: Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the security key field. The firmware works with the ASCII character set only.

The following figure shows the security key entered and confirmed on this dialog.

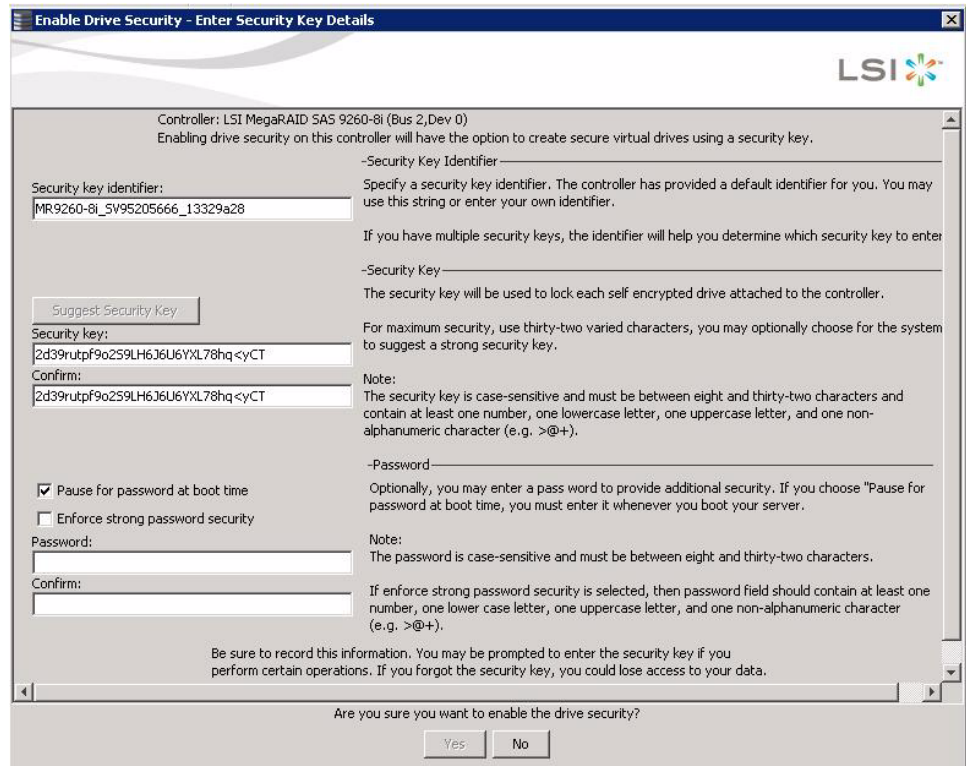


Figure 334: Enable Drive Security – Enter Security Key Details Dialogue

6. (Optional) Select the **Pause for password at boot time** check box.

If you choose this option, you must enter the password whenever you boot the server.

7. (Optional) Select the **Enforce strong password security** check box.

If you choose this option, make sure the password is between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted. The password is case-sensitive.

8. (Optional) Enter a password in the **Password** field and then enter the same password in the **Confirm** field, as shown in the following figure.

Warning messages appear if a mismatch exists between the characters entered in the **Password** field and the **Confirm** field, or if there is an invalid character entered.

NOTE: Be sure to record the password. If you lose the password, you could lose access to your data.

The following figure shows the password entered and confirmed on this dialog.

Enable Drive Security - Enter Security Key Details

Enabling drive security on this controller will have the option to create secure virtual drives using a security key.

Security key identifier:
MR9260-8i_SV95205666_13329a28

Security key:
lj6qeCpUfKRwk2ghw49kezDoq96kjbSI

Confirm:
lj6qeCpUfKRwk2ghw49kezDoq96kjbSI

Pause for password at boot time
 Enforce strong password security

Password:

Confirm:

recorded the security settings for future reference.

Are you sure you want to enable the drive security?
Yes No

Figure 335: Enable Drive Security - Enter Security Key Details Dialog - Password

9. Click **Next**.

The Confirm Enable Drive Security window appears, as shown in the following figure, to show the changes requested to the drive security settings.

NOTE: If you forget the security key, you will lose access to your data. Be sure to record your security key. You might need to enter the security key to perform certain operations.



Figure 336: Confirm Enable Drive Security Dialog

10. Select the **I recorded the security settings for future reference** check box and click **Yes** to confirm that you want to enable drive security on this controller and have recorded the security settings for future reference.

The MegaRAID Storage Manager software enables drive security and returns you to the main menu.

11.6.10 Changing the Drive Security Settings

Perform the following steps to change the encryption settings for the security key identifier, security key, and password.

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Select **Go To >> Controller >> Change Drive Security**.

The Change Security Settings – Introduction dialog appears. This dialog lists the actions you can perform, which include editing the security key identifier, security key, and the password.

3. Click **Next**.

The Change Security Settings - Security Key ID dialog appears.

4. Either keep the existing security key identifier, or enter a new security key identifier.

NOTE: If you change the security key, LSI highly recommends that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

5. Click **Next**.

The Change Security Settings - Security Key dialog appears.

6. Either select the **Use the existing drive security key** radio button to use the existing drive security key, or enter a new security key and then enter the new security key again to confirm.

CAUTION: If you forget the security key, you will lose access to your data. Be sure to record your security key information. You might need to enter the security key to perform certain operations.

The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). The space character is not permitted.

NOTE: Non-U.S. keyboard users must be careful not to enter DBCS characters in the Security Key field. The firmware works with the ASCII character set only.

7. Click **Next**.

The Authenticate Drive Security Settings dialog appears. Authentication is required for the changes that you requested to the drive security settings.

8. Enter the current security key to authenticate the changes.

The Change Security Settings - Password dialog appears.

9. If you choose to, click the option to use a password in addition to the security key.

10. If desired, use a password, either enter the existing password or enter a new password, and enter the password again to confirm.

The text box for the password can hold up to 32 characters. The key must be at least 8 characters.

The next dialog that appears describes the changes you made and asks you whether you want to confirm these changes.

11. Click the check box to confirm that you have recorded the security settings for future reference, and click **Yes** to confirm that you want to change the drive security settings.

The MegaRAID Storage Manager software updates the existing configuration on the controller to use the new security settings and returns you to the main menu.

11.6.11 Disabling Drive Security

NOTE: If you disable drive security, your existing data is not secure and you cannot create any new secure virtual drives. Disabling drive security does not affect the security of data on foreign drives. If you removed any drives that were previously secured, you still need to enter the password when you import them. Otherwise, you cannot access the data on those drives.

NOTE: If any secure drive groups on the controller, you cannot disable drive security. A warning dialog appears if you attempt to do so. To disable drive security, you must first delete the virtual drives on all of the secure drive groups.

Perform the following steps to disable drive security:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
2. Select **Go To >> Controller >> Disable Drive Security**, as shown in the following figure.

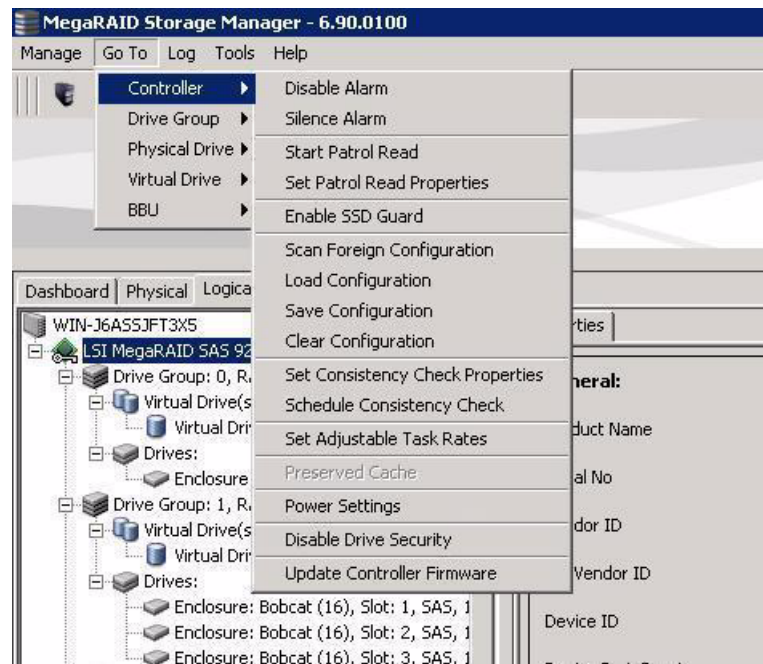


Figure 337: Disable Drive Security Menu Option

The Confirm Disable Drive Security dialog appears.

3. To disable drive security, click **Yes**.

The MegaRAID Storage Manager software disables drive security and returns you to the main menu.

NOTE: If you disable drive security, you cannot create any new encrypted virtual drives and the data on all encrypted unconfigured drives will be erased. Disabling drive security does not affect the security or data of foreign drives.

11.6.12 Importing or Clearing a Foreign Configuration

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the MegaRAID Storage Manager software to import the foreign configuration to the RAID controller or to clear the foreign configuration so you can create a new configuration using these drives.

To import a foreign configuration, you must perform the following tasks:

- Enable security to allow importation of locked foreign configurations. (You can import unsecured or unlocked configurations when security is disabled.)
- Run a scan for foreign configurations.
- If a locked foreign configuration is present and security is enabled, enter the security key, and unlock the configuration.
- Import the foreign configuration.

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.

Verify whether any drives are left to import because the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all the drives are imported, there is no configuration to import.

NOTE: When you create a new configuration, the MegaRAID Storage Manager software shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do **not** appear. To use drives with existing configurations, you must first clear the configuration on those drives.

Perform the following steps to import or clear a configuration:

1. Enable drive security to allow importation of locked foreign drives. See [Section 11.6.9, Enabling Drive Security using LKM](#) for the procedure.
2. After you create a security key, right-click the controller icon, and select **Scan for Foreign Configuration**.

If locked drives (security is enabled) exist, the Unlock Foreign Drives dialog appears.

3. Enter the security key to unlock the configuration.

The Foreign Configuration Detected dialog appears, as shown in the following figure.



Figure 338: Foreign Configuration Detected Dialog

4. Choose one of the following options:
 - Click **Import** to import the foreign configuration from all of the foreign drives.
 - Click **Clear** to remove the configuration from all foreign drives.
 - Click **Advanced** to preview and import specific foreign configurations.
5. Click **OK**.

NOTE: The operation cannot be reversed after it is started. Imported drives display as *Online* in the MegaRAID Storage Manager window.

6. Repeat the import process for any remaining drives.

Because locked drives can use different security key, you must verify whether there are any remaining drives to be imported.

NOTE: When you create a new configuration, the MegaRAID Storage Manager software shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you must first clear the configuration on those drives.

11.6.12.1 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.

The following scenarios can occur with cable pulls or drive removals. Use the Foreign Configuration Preview dialog to import or clear the foreign configuration in each case.

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

- 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 10.2, Running a Consistency Check](#) for more information about checking data consistency.

- 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 10.2, Running a Consistency Check](#) for more information about checking data consistency.

- 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 will be automatically rebuilt. Automatic rebuilds will occur in redundant virtual drives.

- Scenario #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.

11.7 Managing Link Speed

The Managing Link Speed feature allows you to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

All phys in a SAS port can have different link speeds or can have the same link speed.

You can select a link speed setting. However, if phys in a SAS port have different link speed settings and if a phy is connected to a drive or an expander, the firmware overrides the link speed setting you have selected and instead uses the common maximum link speed among all the phys.

To change the link speed, perform the following steps:

1. Right-click a controller in the left frame of the MegaRAID Storage Manager main menu and select **Manage Link Speed**.

or

Select a controller in the left frame of the MegaRAID Storage Manager main menu, and then select **Go To >> Controller >> Manage Link Speed** in the menu bar.

The **Manage Link Speed** dialog box appears, as shown in the following figure.



Figure 339: Manage Link Speed Dialog

- The SAS Address column displays the SAS address that uniquely identifies a device in the SAS domain.
- The Phy column displays the system-supported phy link values. The phy link values are from 0 through 7.
- The Select Link Speed column displays the phy link speeds.

2. Select the desired link speed from the Select Link Speed field using the drop-down selector.

The link speed values are Auto, 1.5, 3.0 or 6.0 Gbps.

NOTE: By default, the link speed in the controller is *Auto* or the value last saved by the user.

3. Click **OK**.

The link speed value is now reset. The change takes place after you restart the system.

The message box appears, as shown in the following figure.

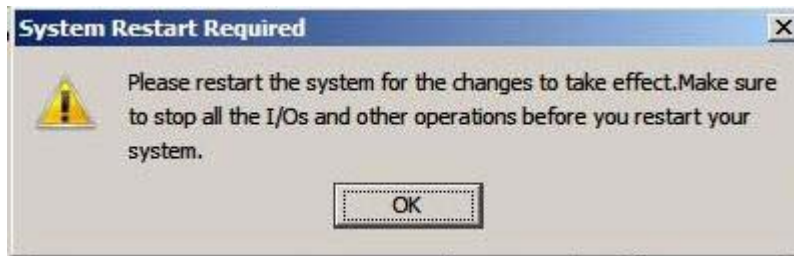


Figure 340: System Restart Required Message

Appendix A

Events and Messages

This appendix lists the MegaRAID Storage Manager events that can appear in the event log.

MegaRAID Storage Manager software monitors the activity and performance of all controllers in the workstation and the devices attached to them. When an event occurs, such as the start of an initialization, an event message appears in the log at the bottom of the MegaRAID Storage Manager window. The messages are also logged in the Windows Application log (Event Viewer).

A.1 Error Levels

Each message that appears in the event log has a Severity level that indicates the severity of the event, as shown in [Table 149](#).

Table 149: Event Error Levels

Severity Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

A.2 Event Messages

[Table 150](#) lists all of the MegaRAID Storage Manager event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. For example, in message No. 1 in the Event Messages table, “%s” is replaced by the firmware version, which is read from the firmware when the event is generated.

Table 150: Event Messages

Number	Severity Level	Event Text
0x0000	Information	MegaRAID firmware initialization started (PCI ID %04x/%04x/%04x/%04x)
0x0001	Information	MegaRAID firmware version %s
0x0002	Fatal	Unable to recover cache data from TBBU
0x0003	Information	Cache data recovered from TBBU successfully
0x0004	Information	Configuration cleared
0x0005	Warning	Cluster down; communication with peer lost
0x0006	Information	Virtual drive %s ownership changed from %02x to %02x
0x0007	Information	Alarm disabled by user

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0008	Information	Alarm enabled by user
0x0009	Information	Background initialization rate changed to %d%%
0x000a	Fatal	Controller cache discarded due to memory/battery problems
0x000b	Fatal	Unable to recover cache data due to configuration mismatch
0x000c	Information	Cache data recovered successfully
0x000d	Fatal	Controller cache discarded due to firmware version incompatibility
0x000e	Information	Consistency Check rate changed to %d%%
0x000f	Fatal	Fatal firmware error: %s
0x0010	Information	Factory defaults restored
0x0011	Information	Flash downloaded image corrupt
0x0012	Critical	Flash erase error
0x0013	Critical	Flash timeout during erase
0x0014	Critical	Flash error
0x0015	Information	Flashing image: %s
0x0016	Information	Flash of new firmware images complete
0x0017	Critical	Flash programming error
0x0018	Critical	Flash timeout during programming
0x0019	Critical	Flash chip type unknown
0x001a	Critical	Flash command set unknown
0x001b	Critical	Flash verify failure
0x001c	Information	Flush rate changed to %d seconds
0x001d	Information	Hibernate command received from host
0x001e	Information	Event log cleared
0x001f	Information	Event log wrapped
0x0020	Fatal	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0021	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0022	Fatal	Not enough controller memory
0x0023	Information	Patrol Read complete
0x0024	Information	Patrol Read paused
0x0025	Information	Patrol Read Rate changed to %d%%
0x0026	Information	Patrol Read resumed
0x0027	Information	Patrol Read started
0x0028	Information	Reconstruction rate changed to %d%%"
0x0029	Information	Drive group modification rate changed to %d%%
0x002a	Information	Shutdown command received from host
0x002b	Information	Test event: %s
0x002c	Information	Time established as %s; (%d seconds since power on)

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x002d	Information	User entered firmware debugger
0x002e	Warning	Background Initialization aborted on %s
0x002f	Warning	Background Initialization corrected medium error (%s at %lx
0x0030	Information	Background Initialization completed on %s
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)
0x0033	Critical	Background Initialization failed on %s
0x0034	Progress	Background Initialization progress on %s is %s
0x0035	Information	Background Initialization started on %s
0x0036	Information	Policy change on %s from %s to %s
0x0038	Warning	Consistency Check aborted on %s
0x0039	Warning	Consistency Check corrected medium error (%s at %lx
0x003a	Information	Consistency Check done on %s
0x003b	Information	Consistency Check done with corrections on %s
0x003c	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)
0x003d	Critical	Consistency Check failed on %s
0x003e	Fatal	Consistency Check completed with uncorrectable data on %s
0x003f	Warning	Consistency Check found inconsistent parity on %s at strip %lx
0x0040	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)
0x0041	Progress	Consistency Check progress on %s is %s
0x0042	Information	Consistency Check started on %s
0x0043	Warning	Initialization aborted on %s
0x0044	Critical	Initialization failed on %s
0x0045	Progress	Initialization progress on %s is %s
0x0046	Information	Fast initialization started on %s
0x0047	Information	Full initialization started on %s
0x0048	Information	Initialization complete on %s
0x0049	Information	LD Properties updated to %s (from %s)
0x004a	Information	Reconstruction complete on %s
0x004b	Fatal	Reconstruction of %s stopped due to unrecoverable errors
0x004c	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)
0x004d	Progress	Reconstruction progress on %s is %s
0x004e	Information	Reconstruction resumed on %s

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x004f	Fatal	Reconstruction resume of %s failed due to configuration mismatch
0x0050	Information	Reconstruction started on %s
0x0051	Information	State change on %s from %s to %s
0x0052	Information	Drive Clear aborted on %s
0x0053	Critical	Drive Clear failed on %s (Error %02x)
0x0054	Progress	Drive Clear progress on %s is %s
0x0055	Information	Drive Clear started on %s
0x0056	Information	Drive Clear completed on %s
0x0057	Warning	Error on %s (Error %02x)
0x0058	Information	Format complete on %s
0x0059	Information	Format started on %s
0x005a	Critical	Hot Spare SMART polling failed on %s (Error %02x)
0x005b	Information	Drive inserted: %s
0x005c	Warning	Drive %s is not supported
0x005d	Warning	Patrol Read corrected medium error on %s at %lx
0x005e	Progress	Patrol Read progress on %s is %s
0x005f	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx
0x0060	Critical	Predictive failure: CDB: %s
0x0061	Fatal	Patrol Read puncturing bad block on %s at %lx
0x0062	Information	Rebuild aborted by user on %s
0x0063	Information	Rebuild complete on %s
0x0064	Information	Rebuild complete on %s
0x0065	Critical	Rebuild failed on %s due to source drive error
0x0066	Critical	Rebuild failed on %s due to target drive error
0x0067	Progress	Rebuild progress on %s is %s
0x0068	Information	Rebuild resumed on %s
0x0069	Information	Rebuild started on %s
0x006a	Information	Rebuild automatically started on %s
0x006b	Critical	Rebuild stopped on %s due to loss of cluster ownership
0x006c	Fatal	Reassign write operation failed on %s at %lx
0x006d	Fatal	Unrecoverable medium error during rebuild on %s at %lx
0x006e	Information	Corrected medium error during recovery on %s at %lx
0x006f	Fatal	Unrecoverable medium error during recovery on %s at %lx
0x0070	Information	Drive removed: %s
0x0071	Warning	Unexpected sense: %s, CDB%s, Sense: %s
0x0072	Information	State change on %s from %s to %s
0x0073	Information	State change by user on %s from %s to %s

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0074	Warning	Redundant path to %s broken
0x0075	Information	Redundant path to %s restored
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group
0x0077	Critical	SAS topology error: Loop detected
0x0078	Critical	SAS topology error: Unaddressable device
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address
0x007a	Critical	SAS topology error: Expander error
0x007b	Critical	SAS topology error: SMP timeout
0x007c	Critical	SAS topology error: Out of route entries
0x007d	Critical	SAS topology error: Index not found
0x007e	Critical	SAS topology error: SMP function failed
0x007f	Critical	SAS topology error: SMP CRC error
0x0080	Critical	SAS topology error: Multiple subtractive
0x0081	Critical	SAS topology error: Table to table
0x0082	Critical	SAS topology error: Multiple paths
0x0083	Fatal	Unable to access device %s
0x0084	Information	Dedicated Hot Spare created on %s (%s)
0x0085	Information	Dedicated Hot Spare %s disabled
0x0086	Critical	Dedicated Hot Spare %s no longer useful for all drive groups
0x0087	Information	Global Hot Spare created on %s (%s)
0x0088	Information	Global Hot Spare %s disabled
0x0089	Critical	Global Hot Spare does not cover all drive groups
0x008a	Information	Created %s}
0x008b	Information	Deleted %s}
0x008c	Information	Marking LD %s inconsistent due to active writes at shutdown
0x008d	Information	Battery Present
0x008e	Warning	Battery Not Present
0x008f	Information	New Battery Detected
0x0090	Information	Battery has been replaced
0x0091	Critical	Battery temperature is high
0x0092	Warning	Battery voltage low
0x0093	Information	Battery started charging
0x0094	Information	Battery is discharging
0x0095	Information	Battery temperature is normal
0x0096	Fatal	Battery has failed and cannot support data retention. Please replace the battery.
0x0097	Information	Battery relearn started
0x0098	Information	Battery relearn in progress

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0099	Information	Battery relearn completed
0x009a	Critical	Battery relearn timed out
0x009b	Information	Battery relearn pending: Battery is under charge
0x009c	Information	Battery relearn postponed
0x009d	Information	Battery relearn will start in 4 days
0x009e	Information	Battery relearn will start in 2 day
0x009f	Information	Battery relearn will start in 1 day
0x00a0	Information	Battery relearn will start in 5 hours
0x00a1	Information	Battery removed
0x00a2	Information	Current capacity of the battery is below threshold
0x00a3	Information	Current capacity of the battery is above threshold
0x00a4	Information	Enclosure (SES) discovered on %s
0x00a5	Information	Enclosure (SAFTE) discovered on %s
0x00a6	Critical	Enclosure %s communication lost
0x00a7	Information	Enclosure %s communication restored
0x00a8	Critical	Enclosure %s fan %d failed
0x00a9	Information	Enclosure %s fan %d inserted
0x00aa	Critical	Enclosure %s fan %d removed
0x00ab	Critical	Enclosure %s power supply %d failed
0x00ac	Information	Enclosure %s power supply %d inserted
0x00ad	Critical	Enclosure %s power supply %d removed
0x00ae	Critical	Enclosure %s SIM %d failed
0x00af	Information	Enclosure %s SIM %d inserted
0x00b0	Critical	Enclosure %s SIM %d removed
0x00b1	Warning	Enclosure %s temperature sensor %d below warning threshold
0x00b2	Critical	Enclosure %s temperature sensor %d below error threshold
0x00b3	Warning	Enclosure %s temperature sensor %d above warning threshold
0x00b4	Critical	Enclosure %s temperature sensor %d above error threshold
0x00b5	Critical	Enclosure %s shutdown
0x00b6	Warning	Enclosure %s not supported; too many enclosures connected to port
0x00b7	Critical	Enclosure %s firmware mismatch
0x00b8	Warning	Enclosure %s sensor %d bad
0x00b9	Critical	Enclosure %s phy %d bad
0x00ba	Critical	Enclosure %s is unstable
0x00bb	Critical	Enclosure %s hardware error
0x00bc	Critical	Enclosure %s not responding

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x00bd	Information	SAS/SATA mixing not supported in enclosure; Drive %s disabled
0x00be	Information	Enclosure (SES) hotplug on %s was detected, but is not supported
0x00bf	Information	Clustering enabled
0x00c0	Information	Clustering disabled
0x00c1	Information	Drive too small to be used for auto-rebuild on %s
0x00c2	Information	BBU enabled; changing WT virtual drives to WB
0x00c3	Warning	BBU disabled; changing WB virtual drives to WT
0x00c4	Warning	Bad block table on drive %s is 80% full
0x00c5	Fatal	Bad block table on drive %s is full; unable to log block %lx
0x00c6	Information	Consistency Check Aborted due to ownership loss on %s
0x00c7	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s
0x00c8	Critical	Battery/charger problems detected; SOH Bad
0x00c9	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded
0x00ca	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded
0x00cb	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled
0x00cc	Critical	Enclosure %s Power supply %d switched off
0x00cd	Information	Enclosure %s Power supply %d switched on
0x00ce	Critical	Enclosure %s Power supply %d cable removed
0x00cf	Information	Enclosure %s Power supply %d cable inserted
0x00d0	Information	Enclosure %s Fan %d returned to normal
0x00d1	Information	BBU Retention test was initiated on previous boot
0x00d2	Information	BBU Retention test passed
0x00d3	Critical	BBU Retention test failed!
0x00d4	Information	NVRAM Retention test was initiated on previous boot
0x00d5	Information	NVRAM Retention test passed
0x00d6	Critical	NVRAM Retention test failed!
0x00d7	Information	%s test completed %d passes successfully
0x00d8	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x
0x00d9	Information	Self check diagnostics completed
0x00da	Information	Foreign Configuration detected
0x00db	Information	Foreign Configuration imported
0x00dc	Information	Foreign Configuration cleared
0x00dd	Warning	NVRAM is corrupt; reinitializing

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x00de	Warning	NVRAM mismatch occurred
0x00df	Warning	SAS wide port %d lost link on PHY %d
0x00e0	Information	SAS wide port %d restored link on PHY %d
0x00e1	Warning	SAS port %d, PHY %d has exceeded the allowed error rate
0x00e2	Warning	Bad block reassigned on %s at %lx to %lx
0x00e3	Information	Controller Hot Plug detected
0x00e4	Warning	Enclosure %s temperature sensor %d differential detected
0x00e5	Information	Drive test cannot start. No qualifying drives found
0x00e6	Information	Time duration provided by host is not sufficient for self check
0x00e7	Information	Marked Missing for %s on drive group %d row %d
0x00e8	Information	Replaced Missing as %s on drive group %d row %d
0x00e9	Information	Enclosure %s Temperature %d returned to normal
0x00ea	Information	Enclosure %s Firmware download in progress
0x00eb	Warning	Enclosure %s Firmware download failed
0x00ec	Warning	%s is not a certified drive
0x00ed	Information	Dirty cache data discarded by user
0x00ee	Information	Drives missing from configuration at boot
0x00ef	Information	Virtual drives (VDs) missing drives and will go offline at boot: %s
0x00f0	Information	VDs missing at boot: %s
0x00f1	Information	Previous configuration completely missing at boot
0x00f2	Information	Battery charge complete
0x00f3	Information	Enclosure %s fan %d speed changed
0x00f4	Information	Dedicated spare %s imported as global due to missing arrays
0x00f5	Information	%s rebuild not possible as SAS/SATA is not supported in an array
0x00f6	Information	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.
0x00f7	Information	Inserted PD: %s Info: %s
0x00f8	Information	Removed PD: %s Info: %s
0x00f9	Information	VD %s is now OPTIMAL
0x00fa	Warning	VD %s is now PARTIALLY DEGRADED
0x00fb	Critical	VD %s is now DEGRADED
0x00fc	Fatal	VD %s is now OFFLINE
0x00fd	Warning	Battery requires reconditioning; please initiate a LEARN cycle
0x00fe	Warning	VD %s disabled because RAID-5 is not supported by this RAID key
0x00ff	Warning	VD %s disabled because RAID-6 is not supported by this controller

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0100	Warning	VD %s disabled because SAS drives are not supported by this RAID key
0x0101	Warning	PD missing: %s
0x0102	Warning	Puncturing of LBAs enabled
0x0103	Warning	Puncturing of LBAs disabled
0x0104	Critical	Enclosure %s EMM %d not installed
0x0105	Information	Package version %s
0x0106	Warning	Global affinity Hot Spare %s commissioned in a different enclosure
0x0107	Warning	Foreign configuration table overflow
0x0108	Warning	Partial foreign configuration imported, PDs not imported:%s
0x0109	Information	Connector %s is active
0x010a	Information	Board Revision %s
0x010b	Warning	Command timeout on PD %s, CDB:%s
0x010c	Warning	PD %s reset (Type %02x)
0x010d	Warning	VD bad block table on %s is 80% full
0x010e	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)
0x010f	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)
0x0110	Information	VD medium error corrected on %s at %lx
0x0111	Warning	Bad block table on PD %s is 100% full
0x0112	Warning	VD bad block table on PD %s is 100% full
0x0113	Fatal	Controller needs replacement, IOP is faulty
0x0114	Information	CopyBack started on PD %s from PD %s
0x0115	Information	CopyBack aborted on PD %s and src is PD %s
0x0116	Information	CopyBack complete on PD %s from PD %s
0x0117	Progress	CopyBack progress on PD %s is %s
0x0118	Information	CopyBack resumed on PD %s from %s
0x0119	Information	CopyBack automatically started on PD %s from %s
0x011a	Critical	CopyBack failed on PD %s due to source %s error
0x011b	Warning	Early Power off warning was unsuccessful
0x011c	Information	BBU FRU is %s
0x011d	Information	%s FRU is %s
0x011e	Information	Controller hardware revision ID %s
0x011f	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata
0x0120	Information	Redundant path restored for PD %s
0x0121	Warning	Redundant path broken for PD %s
0x0122	Information	Redundant enclosure EMM %s inserted for EMM %s

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0123	Information	Redundant enclosure EMM %s removed for EMM %s
0x0124	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD
0x0125	Information	Copyback aborted by user on PD %s and src is PD %s
0x0126	Critical	Copyback aborted on hot spare %s from %s, as hot spare needed for rebuild
0x0127	Warning	Copyback aborted on PD %s from PD %s, as rebuild required in the array
0x0128	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.
0x0129	Information	Copyback cannot be started as PD %s is too small for src PD %s
0x012a	Information	Copyback cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array
0x012b	Information	Microcode update started on PD %s
0x012c	Information	Microcode update completed on PD %s
0x012d	Warning	Microcode update timeout on PD %s
0x012e	Warning	Microcode update failed on PD %s
0x012f	Information	Controller properties changed
0x0130	Information	Patrol Read properties changed
0x0131	Information	CC Schedule properties changed
0x0132	Information	Battery properties changed
0x0133	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled
0x0134	Information	Drive security key created
0x0135	Information	Drive security key backed up
0x0136	Information	Drive security key from escrow, verified
0x0137	Information	Drive security key changed
0x0138	Warning	Drive security key, re-key operation failed
0x0139	Warning	Drive security key is invalid
0x013a	Information	Drive security key destroyed
0x013b	Warning	Drive security key from escrow is invalid
0x013c	Information	VD %s is now secured
0x013d	Warning	VD %s is partially secured
0x013e	Information	PD %s security activated
0x013f	Information	PD %s security disabled
0x0140	Information	PD %s is reprovisioned
0x0141	Information	PD %s security key changed
0x0142	Fatal	Security subsystem problems detected for PD %s

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0143	Fatal	Controller cache pinned for missing or offline VD %s
0x0144	Fatal	Controller cache pinned for missing or offline VDs: %s
0x0145	Information	Controller cache discarded by user for VDs: %s
0x0146	Information	Controller cache destaged for VD %s
0x0147	Warning	Consistency Check started on an inconsistent VD %s
0x0148	Warning	Drive security key failure, cannot access secured configuration
0x0149	Warning	Drive security password from user is invalid
0x014a	Warning	Detected error with the remote battery connector cable
0x014b	Information	Power state change on PD %s from %s to %s
0x014c	Information	Enclosure %s element (SES code 0x%x) status changed
0x014d	Information	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group
0x014e	Information	Copyback cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group
0x014f	Information	VD bad block table on %s is cleared
0x0150	Caution	SAS topology error: 0x%x
0x0151	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)
0x0152	Information	Controller requests a host bus rescan
0x0153	Information	Controller repurposed and factory defaults restored
0x0154	Information	Drive security key binding updated
0x0155	Information	Drive security is in EKM mode
0x0156	Warning	Drive security failed to communicate with EKMS
0x0157	Information	%s needs key to be %s %s
0x0158	Warning	%s secure failed
0x0159	Critical	Controller encountered a fatal error and was reset
0x015a	Information	Snapshots enabled on %s (Repository %s)
0x015b	Information	Snapshots disabled on %s (Repository %s) by the user
0x015c	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error
0x015d	Information	Snapshot created on %s at %s
0x015e	Information	Snapshot deleted on %s at %s
0x015f	Information	View created at %s to a snapshot at %s for %s
0x0160	Information	View at %s is deleted, to snapshot at %s for %s
0x0161	Information	Snapshot rollback started on %s from snapshot at %s
0x0162	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s
0x0163	Information	Snapshot rollback on %s completed for snapshot at %s
0x0164	Information	Snapshot rollback progress for snapshot at %s, on %s is %s
0x0165	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0166	Critical	Snapshot space for %s in snapshot repository %s, is full
0x0167	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s
0x0168	Critical	View at %s to snapshot at %s, is full on snapshot repository %s
0x0169	Critical	Snapshot repository lost for %s
0x016a	Warning	Snapshot repository restored for %s
0x016b	Critical	Snapshot encountered an unexpected internal error: 0x%lx
0x016c	Information	Auto Snapshot enabled on %s (snapshot repository %s)
0x016d	Information	Auto Snapshot disabled on %s (snapshot repository %s)
0x016e	Critical	Configuration command could not be committed to disk, please retry
0x016f	Information	COD on %s updated as it was stale
0x0170	Warning	Power state change failed on %s (from %s to %s)
0x0171	Warning	%s is not available
0x0172	Information	%s is available
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks
0x0174	Information	%s is using CacheCade %s
0x0175	Information	%s is no longer using CacheCade %s
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s
0x0177	Warning	Auto Snapshot failed for %s in snapshot repository %s
0x0178	Warning	Controller reset on-board expander
0x0179	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks
0x017a	Warning	Battery cannot initiate transparent learn cycles
0x017b	Information	Premium feature %s key was applied for - %s
0x017c	Information	Snapshot schedule properties changed on %s
0x017d	Information	Snapshot scheduled action is due on %s
0x017e	Information	Performance Metrics: collection command 0x%lx
0x017f	Information	Premium feature %s key was transferred - %s
0x0180	Information	Premium feature serial number %s
0x0181	Warning	Premium feature serial number mismatched. Key-vault serial num - %s
0x0182	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery
0x0183	Information	%s power policy changed to %s (from %s)
0x0184	Warning	%s cannot transition to max power savings
0x0185	Information	Host driver is loaded and operational
0x0186	Information	%s mirror broken
0x0187	Information	%s mirror joined

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x0188	Warning	%s link %d failure in wide port
0x0189	Information	%s link %d restored in wide port
0x018a	Information	Memory module FRU is %s
0x018b	Warning	Cache-vault power pack is sub-optimal. Please replace the pack
0x018c	Warning	Foreign configuration auto-import did not import any drives
0x018d	Warning	Cache-vault microcode update required
0x018e	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used
0x018f	Warning	LD (%s) protection information lost
0x0190	Information	Diagnostics passed for %s
0x0191	Critical	Diagnostics failed for %s
0x0192	Information	Server Power capability Diagnostic Test Started
0x0193	Information	Drive Cache settings enabled during rebuild for %s
0x0194	Information	Drive Cache settings restored after rebuild for %s
0x0195	Information	Drive %s commissioned as Emergency spare
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare
0x0197	Information	Consistency Check suspended on %s
0x0198	Information	Consistency Check resumed on %s
0x0199	Information	Background Initialization suspended on %s
0x019a	Information	Background Initialization resumed on %
0x019b	Information	Reconstruction suspended on %s
0x019c	Information	Rebuild suspended on %
0x019d	Information	Copyback suspended on %s
0x019e	Information	Reminder: Consistency Check suspended on %
0x019f	Information	Reminder: Background Initialization suspended on %s
0x01a0	Information	Reminder: Reconstruction suspended on %s
0x01a1	Information	Reminder: Rebuild suspended on %s
0x01a2	Information	Reminder: Copyback suspended on %s
0x01a3	Information	Reminder: Patrol Read suspended
0x01a4	Information	Erase aborted on %s
0x01a5	Critical	Erase failed on %s (Error %02x)
0x01a6	Progress	Erase progress on %s is %s
0x01a7	Information	Erase started on %s
0x01a8	Information	Erase completed on %s
0x01a9	Information	Erase aborted on %s
0x01aa	Critical	Erase failed on %s
0x01ab	Progress	Erase progress on %s is %s

Table 150: Event Messages (Continued)

Number	Severity Level	Event Text
0x01ac	Information	Erase started on %s
0x01ad	Information	Erase complete on %s
0x01ae	Warning	Potential leakage during erase on %s
0x01af	Warning	Battery charging was suspended due to high battery temperature
0x01b0	Information	NVCache firmware update was successful
0x01b1	Warning	NVCache firmware update failed
0x01b2	Fatal	%s access blocked as cached data in CacheCade is unavailable
0x01b3	Information	CacheCade disassociate started on %s
0x01b4	Information	CacheCade disassociate completed on %s
0x01b5	Critical	CacheCade disassociate failed on %s
0x01b6	Progress	CacheCade disassociate progress on %s is %s
0x01b7	Information	CacheCade disassociate aborted by user on %s
0x01b8	Information	Link speed changed on SAS port %d and PHY %d
0x01b9	Warning	Advanced Software Options was deactivated for - %s
0x01ba	Information	%s is now accessible
0x01bb	Information	%s is using CacheCade
0x01bc	Information	%s is no longer using CacheCade
0x01bd	Information	Patrol Read aborted on %s

Appendix B

MegaCLI Error Messages

This appendix lists the MegaCLI error messages.

The MegaCLI Configuration Utility is a command line interface application you can use to manage MegaRAID SAS RAID controllers. See [Chapter 5](#) for more information about the MegaCLI utility and commands.

B.1 Error Messages and Descriptions

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in [Table 151](#).

Table 151: Error Messages and Descriptions

Number	Event Text
0x00	Command completed successfully
0x01	Invalid command
0x02	DCMD opcode is invalid
0x03	Input parameters are invalid
0x04	Invalid sequence number
0x05	Abort isn't possible for the requested command
0x06	Application 'host' code not found
0x07	Application already in use - try later
0x08	Application not initialized
0x09	Given array index is invalid
0x0a	Unable to add missing drive to array, as row has no empty slots
0x0b	Some of the CFG resources conflict with each other or the current config
0x0c	Invalid device ID / select-timeout
0x0d	Drive is too small for requested operation
0x0e	Flash memory allocation failed
0x0f	Flash download already in progress
0x10	Flash operation failed
0x11	Flash image was bad
0x12	Downloaded flash image is incomplete
0x13	Flash OPEN was not done
0x14	Flash sequence is not active
0x15	Flush command failed
0x16	Specified application doesn't have host-resident code

Table 151: Error Messages and Descriptions (Continued)

Number	Event Text
0x17	LD operation not possible - CC is in progress
0x18	LD initialization in progress
0x19	LBA is out of range
0x1a	Maximum LDs are already configured
0x1b	LD is not OPTIMAL
0x1c	LD Rebuild is in progress
0x1d	LD is undergoing reconstruction
0x1e	LD RAID level is wrong for requested operation
0x1f	Too many spares assigned
0x20	Scratch memory not available - try command again later
0x21	Error writing MFC data to SEEPROM
0x22	Required HW is missing (i.e. Alarm or BBU)
0x23	Item not found
0x24	LD drives are not within an enclosure
0x25	PD CLEAR operation is in progress
0x26	Unable to use SATA(SAS) drive to replace SAS(SATA)
0x27	Patrol Read is disabled
0x28	Given row index is invalid
0x2d	SCSI command done, but non-GOOD status was received-see mf.hdr.extStatus for SCSI_STATUS
0x2e	IO request for MFI_CMD_OP_PD_SCSI failed - see extStatus for DM error
0x2f	Matches SCSI RESERVATION_CONFLICT
0x30	One or more of the flush operations failed
0x31	Firmware real-time currently not set
0x32	Command issues while firmware in wrong state (i.e., GET RECON when op not active)
0x33	LD is not OFFLINE - IO not possible
0x34	Peer controller rejected request (possibly due to resource conflict)
0x35	Unable to inform peer of communication changes (retry might be appropriate)
0x36	LD reservation already in progress
0x37	I2C errors were detected
0x38	PCI errors occurred during XOR/DMA operation
0x39	Diagnostics failed - see event log for details
0x3a	Unable to process command as boot messages are pending
0x3b	Returned in case if foreign configurations are incomplete
0x3d	Returned in case if a command is tried on unsupported hardware

Table 151: Error Messages and Descriptions (Continued)

Number	Event Text
0x3e	CC scheduling is disabled
0x3f	PD CopyBack operation is in progress
0x40	Selected more than one PD per array
0x41	Microcode update operation failed
0x42	Unable to process command as drive security feature is not enabled
0x43	Controller already has a lock key
0x44	Lock key cannot be backed-up
0x45	Lock key backup cannot be verified
0x46	Lock key from backup failed verification
0x47	Rekey operation not allowed, unless controller already has a lock key
0x48	Lock key is not valid, cannot authenticate
0x49	Lock key from escrow cannot be used
0x4a	Lock key backup (pass-phrase) is required
0x4b	Secure LD exist
0x4c	LD secure operation is not allowed
0x4d	Reprovisioning is not allowed
0x4e	Drive security type (FDE or non-FDE) is not appropriate for requested operation
0x4f	LD encryption type is not supported
0x50	Cannot mix FDE and non-FDE drives in same array
0x51	Cannot mix secure and unsecured LD in same array
0x52	Secret key not allowed
0x53	Physical device errors were detected
0x54	Controller has LD cache pinned
0x55	Requested operation is already in progress
0x56	Another power state set operation is in progress
0x57	Power state of device is not correct
0x58	No PD is available for patrol read
0x59	Controller reset is required
0x5a	No EKM boot agent detected
0x5b	No space on the snapshot repository VD
0x5c	For consistency SET PiTs, some PiT creations might fail and some succeed
0xFF	Invalid status - used for polling command completion

Appendix C

Glossary

This appendix provides a glossary for terms used in this document.

access policy	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> .
alarm enabled	A controller property that indicates whether the controller's onboard alarm is enabled.
alarm present	A controller property that indicates whether the controller has an onboard alarm. If present and enabled, the alarm is sounded for certain error conditions.
array	See <i>drive group</i> .
BBU present	A controller property that indicates whether the controller has an onboard battery backup unit to provide power in case of a power failure.
BGI rate	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
BIOS	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.
cache	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.
cache flush interval	A controller property that indicates how often the data cache is flushed.
caching	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 can temporarily store data in accordance with its write back policies.
capacity	A property that indicates the amount of storage space on a drive or virtual drive.
coerced capacity	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-GB drive from one manufacturer might be 4,196 MB, and a 4-GB from another manufacturer might be 4,128 MB. These drives could be coerced to a usable capacity of 4,088 MB each for use in a drive group in a storage configuration.
coercion mode	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.

consistency check	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.
consistency check rate	The rate at which consistency check operations are run on a computer system.
controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection. MegaRAID Storage Manager software runs on LSI SAS controllers.
copyback	<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>
current write policy	<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"> ■ 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.
default write policy	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.
device ID	A controller or drive property indicating the manufacturer-assigned device ID.
device port count	A controller property indicating the number of ports on the controller.
drive cache policy	A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.
drive group	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.
drive state	<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"> ■ Unconfigured Good – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. ■ Hot Spare – A drive that is configured as a hot spare. ■ Online – A drive that can be accessed by the RAID controller and will be part of the virtual drive. ■ Rebuild – A drive to which data is being written to restore full redundancy for a virtual drive. ■ Failed – A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.

	<ul style="list-style-type: none"> ■ 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. ■ None – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.
drive subsystem	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.
drive type	A drive property indicating the characteristics of the drive.
EKM	External Key Management
fast initialization	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.
fault tolerance	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.
firmware	Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a system that loads the full operating system from drive or from a network and then passes control to the operating system.
foreign configuration	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.
formatting	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.
hole	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.
host interface	A controller property indicating the type of interface used by the computer host system: for example, <i>PCIX</i> .
host port count	A controller property indicating the number of host data ports currently in use.
host system	Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.
hot spare	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. 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.

initialization	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.
IO policy	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.)
learning cycle	A battery calibration operation performed by a RAID controller periodically to determine the condition of the battery.
load-balancing	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.
LKM	Local Key Management
media error count	A drive property indicating the number of errors that have been detected on the drive media.
migration	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.
mirroring	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.
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.
name	A virtual drive property indicating the user-assigned name of the virtual drive.
non-redundant configuration	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.
NVRAM	Acronym for non-volatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.
NVRAM present	A controller property indicating whether an NVRAM is present on the controller.
NVRAM size	A controller property indicating the capacity of the controller's NVRAM.
offline	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.

patrol read	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 might not be necessary.
patrol read rate	The user-defined rate at which patrol read operations are run on a computer system.
product info	A drive property indicating the vendor-assigned model number of the drive.
product name	A controller property indicating the manufacturing name of the controller.
RAID	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.
RAID 0	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.
RAID 00	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.
RAID 1	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.
RAID 5	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	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.
RAID 10	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.
RAID 50	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.
RAID 60	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.
RAID level	A virtual drive property indicating the RAID level of the virtual drive. LSI SAS RAID controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.

read policy	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 (known as Normal mode in WebBIOS), read ahead capability is disabled.
rebuild	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.
rebuild rate	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.
reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click Reclaim , the individual drives are removed from the virtual drive configuration.
reconstruction rate	The user-defined rate at which a drive group modification operation is carried out.
redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
redundant configuration	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.
reversible hot spare	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.
revision level	A drive property that indicates the revision level of the drive's firmware.
SAS	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.
SATA	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.
SCSI device type	A drive property indicating the type of the device, such as drive.
serial no.	A controller property indicating the manufacturer-assigned serial number.
strip size	The portion of a stripe that resides on a single drive in the drive group.
stripe size	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.

striping	<p>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.</p>
subvendor ID	<p>A controller property that lists additional vendor ID information about the controller.</p>
uncorrectable error count	<p>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.</p>
vendor ID	<p>A controller property indicating the vendor-assigned ID number of the controller.</p>
vendor info	<p>A drive property listing the name of the vendor of the drive.</p>
virtual drive	<p>A storage unit created by a RAID controller from one or more drives. Although a virtual drive can 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 can retain redundant data in case of a drive failure.</p>
virtual drive state	<p>A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.</p>
write-back	<p>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.</p> <p>These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.</p>
write policy	<p>See <i>Default Write Policy</i>.</p>
write-through	<p>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.</p>

Appendix D

History of Technical Changes

This appendix lists all the technical changes made to this guide for all the previous releases.

Table 152: History of Technical Changes

Version and Date	Description of Changes
80-00156-01 Rev. K, February 2011	<p>Added Shield State on page 152 in the WebBIOS section</p> <p>Added Viewing and Changing Battery Backup Unit Information on page 155 in the WebBIOS section.</p> <p>Added Viewing Enclosure Properties on page 162 in the WebBIOS section.</p> <p>Added SSD Disk Cache Policy on page 167 in the WebBIOS section.</p> <p>Added Emergency Hotspare on page 167 in the WebBIOS section.</p> <p>Added Emergency Hotspare for Controllers on page 168 in the WebBIOS section.</p> <p>Updated Viewing and Expanding a Virtual Drive on page 172 in the WebBIOS section.</p> <p>Added Shield State on page 314 in the MSM section.</p> <p>Added Logical View Shield State on page 315 in the MSM section.</p> <p>Added Viewing the Physical Drive Properties on page 315 in the MSM section.</p> <p>Added Viewing Server Profile of a Drive in Shield State on page 316 in the MSM section.</p> <p>Added Displaying the Virtual Drive Properties on page 317 in the MSM section.</p> <p>Added Emergency HotSpare Support on page 320 in the MSM section.</p> <p>Added SSD Disk Cache Policy on page 322 in the MSM section.</p> <p>Added Non-SED Secure Erase Support on page 324 in the MSM section.</p> <p>Added Rebuild Write Cache on page 329 in the MSM section.</p> <p>Added Background Suspend or Resume Support on page 330 in the MSM section.</p> <p>Added Enclosure Properties on page 331 in the MSM section.</p> <p>Updated Monitoring Battery Backup Units on page 331 in the MSM section.</p>
80-00156-01, Rev. J, September 2010	<p>Added Managing Software Licensing on page 56 in the WebBIOS section.</p> <p>Added EKM and LKM on page 127 in the WebBIOS section.</p> <p>Added Import Foreign Drives in EKM/EKM Secured Locked Drives on page 204.</p> <p>Added Enable the Snapshot Scheduler on page 269 in the WebBIOS section.</p> <p>Added WebBIOS Dimmer Switch on page 208 in the WebBIOS section.</p> <p>Added SafeStore Security Options on page 227 details in the MegaCLI section.</p> <p>Added section Enable the Snapshot Scheduler on page 269 in the MSM section.</p> <p>Added MegaRAID Software Licensing on page 407 in the MSM section.</p> <p>Added Enabling Drive Security using EKM on page 452 in the MSM section.</p> <p>Added Enhanced Dimmer Switch Power Settings on page 356 in the MSM section.</p> <p>Added Software Licensing, EKM and LKM, Dimmer Switch, Software Licensing, and other sections in Chapter 5.</p>

Table 152: History of Technical Changes

Version and Date	Description of Changes
80-00156-01, Rev. I, June 2010	Updated the document with changes to the software utilities. Added Chapter 11 for the MegaRAID advanced software features.
80-00156-01, Rev. H, July 2009	Documented the Full Disk Encryption (FDE) feature.
80-00156-01, Rev. G, June 2009	Updated the MegaRAID Storage Manager chapters.
80-00156-01, Rev. F, March 2009	Updated the WebBIOS Configuration Utility, MegaRAID Storage Manager, and MegaCLI chapters.
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.	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.



51530-00B

